arXiv:1702.03917v1 [math.PR] 13 Feb 2017

MEXIT: Maximal un-coupling times for Markov processes

Philip A. Ernst*, Wilfrid S. Kendall†, Gareth O. Roberts‡, Jeffrey S. Rosenthal§

14th February 2017

Abstract

Classical coupling constructions arrange for copies of the *same* Markov process started at two *different* initial states to become equal as soon as possible. In this paper, we consider an alternative coupling framework in which one seeks to arrange for two *different* Markov processes to remain equal for as long as possible, when started in the *same* state. We refer to this "un-coupling" or "maximal agreement" construction as *MEXIT*, standing for "maximal exit" time. After highlighting the importance of un-coupling arguments in a few key statistical and probabilistic settings, we develop an explicit *MEXIT* construction for Markov chains in discrete time with countable state-space. This construction is generalized to random processes on general state-space running in continuous time, and then exemplified by discussion of *MEXIT* for Brownian motions with two different constant drifts.

MSC 2010: 60J60; 60G05

Keywords: ADAPTIVE MCMC; COPULA; COUPLING; DIFFUSIONS; FRÉCHET CLASS; HAHN-JORDAN DECOMPOSITION; MARKOVIAN COUPLING; MCMC; MEET MEASURE; MEXIT; ONE-STEP MINORIZATION; PSEUDO-MARGINAL MCMC; RECOGNITION LEMMA FOR MAXIMAL COUPLING; UN-COUPLING.

1 Introduction

Coupling is a device commonly employed in probability theory for learning about distributions of certain random variables by means of judicious construction in ways which depend on other random variables (Lindvall, 1992; Thorisson, 2000). Such coupling constructions are often used to prove convergence of Markov processes to stationary distributions (Pitman, 1976), especially for Markov chain Monte Carlo (MCMC) algorithms (Roberts and Rosenthal, 2004, and references therein), by seeking to build two different copies of the *same* Markov process started at two *different* initial states in such a way that they become equal at a fast rate. Fastest possible rates are achieved by the *maximal coupling* constructions which were introduced and studied in Griffeath (1975), Pitman (1976), and Goldstein (1979). Our results and methods are closely related to the work of Goldstein (1979), who deals with rather general discrete-time random processes. Our situation is related to a time-reversal of the situation studied by Goldstein. However we are able to avoid some measure-theoretic assumptions, since we work initially with finite time intervals and focus on explicit construction of probability measures.

In the current work, we consider what might be viewed as the dual problem where coupling is used to try to construct two *different* Markov processes which remain equal for as long as possible, when they are started in the *same* state. That is, we move from consideration of the coupling time to focus on the *un-coupling time* at which the processes diverge, and try to make that as *large* as possible. We refer to this as *MEXIT* (standing for "maximal exit" time). While finalizing our current work, it came to our attention that this construction is the same as the *maximal agreement coupling time* of the August 2016 work of Völlering, who additionally derives a lower bound on *MEXIT*. Nonetheless, we believe the current work complements Völlering well. It offers an explicit treatment of discrete-time countable-state-space, generalizes the continuous-time case, gives an explicit description of the *MEXIT* construction, and discusses a number of significant applications of *MEXIT*. We note that the work of Völlering (2016) does not consider the continuous-time case.

In addition to being a natural mathematical question, *MEXIT* has direct applications to many key statistical and probabilistic settings (see Section 2 below). In particular, couplings which are *Markovian* and *faithful*

*philip.ernst@rice.edu
†w.s.kendall@warwick.ac.uk
‡gareth.o.roberts@warwick.ac.uk
§jeff@math.toronto.edu

(Rosenthal, 1997; alternatively "co-adapted" or "immersion", depending on the extent to which one wishes to emphasize the underlying filtration as in Burdzy and Kendall, 2000; Kendall, 2015) are the most straightforward to construct, but often are *not* maximal, while more complicated non-Markovian and non-faithful couplings lead to stronger bounds. The same is true in the context of *MEXIT*.

2 Applications

To motivate the natural role of *MEXIT* in the existing literature, we first consider the role of un-coupling arguments in a few statistical and probabilistic settings.

2.1 Bounds on accuracy for statistical tests

Un-coupling has an impact on the theory classical statistical testing. In Farrell (1964), amongst other sources, some function of the data (but not the data itself) is assumed to have been observed. A statistical test is then constructed to enable detection of the distribution from which the observed data have been sampled. For example, suppose that data $X_1, X_2, \ldots$ are generated as a draw either from a multivariate probability distribution $\mathbb{P}_1$ or from a multivariate probability distribution $\mathbb{P}_2$. The goal is to determine whether the data was generated from $\mathbb{P}_1$ or from $\mathbb{P}_2$. For some function h of the data, and some acceptance region A, the statistical test decides in favor of $\mathbb{P}_1$ if $h(X_1, \ldots, X_n) \in A$ and otherwise decides in favor of $\mathbb{P}_2$.

Suppose that there exists an un-coupling time T, such that if $X_1, X_2, \ldots$ are generated from $\mathbb{P}_1$, and if $Y_1, Y_2, \ldots$ are generated from $\mathbb{P}_2$ then it is exactly the case that $X_i = Y_i$ for all $1 \leq i \leq T$ (so that $X_i \neq Y_i$ for all $i > T$). We use $\mathbb{P}$ to refer to the joint distribution (in fact, the coupling) of $\mathbb{P}_1$ and $\mathbb{P}_2$.

The following proposition uses the un-coupling probabilities to recover a lower bound on the accuracy of such statistical tests related to Farrell (1964, Theorem 1).

Proposition 1. *Under the above assumptions, the sum of the probabilities of Type-I and Type-II errors of our statistical test is at least* $\mathbb{P}[T > n]$.

Proof. We apply elementary arguments to the sum of the probabilities of Type-I and Type-II errors:

$$
\begin{aligned}
&\mathbb{P}_2[h(Y_1, \ldots, Y_n) \in A] + \mathbb{P}_1[h(X_1, \ldots, X_n) \notin A] \quad = \\
&\quad = \quad \mathbb{P}_2[h(Y_1, \ldots, Y_n) \in A] + 1 - \mathbb{P}_1[h(X_1, \ldots, X_n) \in A] \\
&\quad = \quad 1 - \Big(\mathbb{P}_1[h(Y_1, \ldots, Y_n) \in A] - \mathbb{P}_2[h(X_1, \ldots, X_n) \in A]\Big) \\
&\quad \geq \quad 1 - |\mathbb{P}_1[h(Y_1, \ldots, Y_n] \in A] - \mathbb{P}_2[h(X_1, \ldots, X_n) \in A]| \\
&\quad \geq \quad 1 - \|\mathcal{L}_{\mathbb{P}_1}(X_1, \ldots, X_n) - \mathcal{L}_{\mathbb{P}_2}(Y_1, \ldots, Y_n)\| \quad \text{(definition of total variation distance)} \\
&\quad \geq \quad 1 - \mathbb{P}[X_i \neq Y_i \text{ for some } 1 \leq i \leq n] \quad \text{(coupling inequality)} \\
&\quad = \quad 1 - (1 - \mathbb{P}[X_i = Y_i \text{ for all } \ 1 \leq i \leq n]) \\
&\quad = \quad \mathbb{P}[X_i = Y_i \text{ for all } 1 \leq i \leq n] \quad = \quad \mathbb{P}[T > n] \ .
\end{aligned}
$$

□

2.2 Two independent coin flips

We now turn to the classical probabilistic paradigm of coin flips. Let X and Y represent two different sequences of i.i.d. coin flips, with probabilities of landing on H (heads) to be q and r respectively, where $0 \leq r \leq q \leq 1/2$. Suppose that we wish to maximise the length of the initial segment for which coin flips agree:

$$
S \quad = \quad \max\{t : X_i = Y_i \text{ for all } 1 \leq i \leq t\} \, .
$$

For concreteness, we will set $q = 0.5$ and $r = 0.4$ throughout this section; the generalization to other values is immediate.

2.2.1 Markovian Faithful Coupling for Independent Coin Flips

The "greedy" (Markovian and faithful) coupling carries out the best "one-step minorization" coupling possible, separately at each iteration. One-step minorization is essentially maximal coupling for single random variables. In this case, that means that for each flip, $\mathbb{P}[X = Y = H] = 0.4$, $\mathbb{P}[X = Y = T] = 0.5$, and $\mathbb{P}[X = H,\ Y = T] = 0.1$. This preserves the marginal distributions of X and Y, and yields $\mathbb{P}[X = Y] = 0.9$ at each step. Accordingly, the probability of agreement continuing for at least n steps is given by $\mathbb{P}[X_i = Y_i \text{ for } 1 \leq i \leq n] = (0.9)^n$.

2.2.2 A Look-ahead Coupling for Independent Coin Flips

Let a "look-ahead" coupling be a coupling which instead uses an n-step minorization couple on the entire sequence of n coin tosses, so that for each sequence s of n different Heads and Tails, it sets $\mathbb{P}[X = Y = s] = \min(\mathbb{P}[X = s], \mathbb{P}[Y = s])$. Consequently, if s has m Heads and $n - m$ Tails, then

$$\mathbb{P}[X = Y = s] \quad = \quad \min\{0.5^n,\, 0.4^m 0.6^{n-m}\}\,.$$

Elementary events for which X and Y disagree are assigned probabilities which preserve the marginal distributions of X and of Y. The simplest way to implement this is to use "independent residuals", but other choices are also possible.

This look-ahead coupling leads to a larger probability that $X = Y$. Indeed, even in the case $n = 2$, the probability of agreement over two coin flips under the greedy coupling is given by

$$\mathbb{P}[X = Y] \quad = \quad (0.9)^2 = 0.81\,.$$

The look-ahead coupling delivers a strictly larger probability of agreement over two coin flips:

$$\begin{aligned}\mathbb{P}[X = Y] \quad &= \quad \min(0.5^2, 0.4^2) + \min(0.5^2, 0.6^2) + 2\min(0.5^2, 0.4 \cdot 0.6) \\ &= \quad 0.4^2 + 0.5^2 + 2 \cdot 0.4 \cdot 0.6 = 0.16 + 0.25 + 0.48 = 0.89\,.\end{aligned}$$

When $n = 2$, the matrix of joint probabilities for X and Y under the look-ahead coupling is calculated to be:

$X \backslash Y$	HH	HT	TH	TT	SUM
HH	0.16	0	0	0.09	0.25
HT	0	0.24	0	0.01	0.25
TH	0	0	0.24	0.01	0.25
TT	0	0	0	0.25	0.25
SUM	0.16	0.24	0.24	0.36	1

Marginalizing this coupling on the initial coin flip ("projecting back" to the initial flip, with $n = 1$), we see that $\mathbb{P}[X_1 = Y_1 = H] = 0.16 + 0.24 = 0.4$, and $\mathbb{P}[X_1 = Y_1 = T] = 0.24 + 0.01 + 0.25 = 0.5$, and $\mathbb{P}[X_1 = H,\ Y_1 = T] = 0.09 + 0.01 = 0.1$. The projection to the initial flip yields the same agreement probability as would have been attained by maximizing the probability of staying together for just one flip ($n = 1$). That is, the $n = 2$ look-ahead coupling construction is *compatible* with the $n = 1$ construction.

Finally, it is worth noting that the $n = 2$ look-ahead coupling is certainly not faithful. For example, $\mathbb{P}[X_2 = H \,|\, X_1 = Y_1 = H] = 0.4 \neq 0.5$, and $\mathbb{P}[X_2 = H \,|\, X_1 = H,\ Y_1 = T] = 0.9 \neq 0.5$, *etc.*

2.3 A Look-ahead coupling for independent coin flips: the case $n = 3$

The matrix of joint probabilities for X and Y under the look-ahead coupling for $n = 3$ is more complicated, but can be calculated as:

$X \backslash Y$	HHH	HHT	HTH	HTT	THH	THT	TTH	TTT	SUM
HHH	0.064	0	0	0.0078	0	0.0078	0.0078	0.0375	0.125
HHT	0	0.096	0	0.0037	0	0.0037	0.0037	0.0178	0.125
HTH	0	0	0.096	0.0037	0	0.0037	0.0037	0.0178	0.125
HTT	0	0	0	0.125	0	0	0	0	0.125
THH	0	0	0	0.0037	0.096	0.0037	0.0037	0.0178	0.125
THT	0	0	0	0	0	0.125	0	0	0.125
TTH	0	0	0	0	0	0	0.125	0	0.125
TTT	0	0	0	0	0	0	0	0.125	0.125
SUM	0.064	0.096	0.096	0.144	0.096	0.144	0.144	0.216	1

With these probabilities, we compute that

$$\mathbb{P}[X = Y] = 0.064 + 3 \times 0.096 + 4 \times 0.125 = 0.852\,.$$

This is greater than the agreement probability of $0.9^3 = 0.729$ that would have be achieved via the greedy coupling. It is natural to wonder whether or not it is possible always to ensure that such a construction works not just for one fixed time but for all times. We further expound on this point in Sections 3 and 4, where discussion of a much more general context shows that that such constructions always exist.

2.3.1 Optimal Expectation

Until now, this section has focused on maximising $\mathbb{P}[X_i = Y_i \text{ for all } 1 \leq i \leq n]$, which is to say, maximizing $\mathbb{P}[S \geq n]$ with S being the time of first disagreement as above. We now consider the related question of maximizing the expected value $\mathbb{E}[S]$. Using the greedy coupling, clearly

$$\mathbb{E}[S] \quad = \quad \sum_{j=1}^{\infty} \mathbb{P}[S \geq j] \quad = \quad \sum_{j=1}^{\infty} 0.9^j \quad = \quad 0.9/(1-0.9) = 9\,.$$

If the different look-ahead couplings are chosen to be compatible, then this shows that $\mathbb{E}[S]$ is the sum for $r = 1, 2, \ldots$ of the probabilities that the j^{th} look-ahead coupling was successful. The work of Sections 3 and 4 shows that such a choice is always feasible, even for very general random processes indeed.

2.4 Adaptive MCMC

Un-coupling arguments play a natural role in the adaptive MCMC (Markov-chain Monte Carlo) literature, highlighted in particular by the work of Roberts and Rosenthal (2007). Roberts and Rosenthal (2007) prove convergence of *adaptive* MCMC by comparing an adaptive process to a process which "*stops* adapting" at some point, and then by showing that the two processes have a high probability of remaining equal long enough such that the second process (and hence also the first process) converge to stationarity. The authors accomplish this by considering a sequence of adaptive Markov kernels $P_{\Gamma_1}, P_{\Gamma_2}, \ldots$ on a state space $\mathcal{X}$, where $\{P_\gamma : \gamma \in \mathcal{Y}\}$ are a collection of Markov kernels each having the same stationary probability distribution π, and the Γ_i are $\mathcal{Y}$-valued random variables which are "adaptive" (i.e., they depend on the previous Markov chain values but not on future values). Under appropriate assumptions, the authors prove that a Markov chain X which evolves *via* the adaptive Markov kernels will still converge to the specified stationary distribution π. The key step in the proof of the central result (Roberts and Rosenthal, 2007, Theorem 5) is an un-coupling approach, highlighted below.

Roberts and Rosenthal (2007, Theorem 5) assume that, for any $\varepsilon > 0$, there is a non-negative integer $N = N(\varepsilon)$ such that

$$\|P_\gamma^N(x, \cdot) - \pi(\cdot)\|_{\text{TV}} \quad \leq \quad \varepsilon$$

for all $x \in \mathcal{X}$ and $\gamma \in \mathcal{Y}$ (where $\|\cdot\|_{\text{TV}}$ denotes total variation norm of a signed measure). Furthermore, there is a non-negative integer $n^* = n^*(\varepsilon)$ such that with probability at least $1 - \varepsilon/N$,

$$\sup_{x \in \mathcal{X}} \|P_{\Gamma_{n+1}}(x, \cdot) - P_{\Gamma_n(x,\cdot)}\|_{\text{TV}} \quad \leq \quad \varepsilon/N^2$$

for all $n \geq n^*$.

These assumptions are used to prove, for any $K \geq n^* + N$, the existence of a pair of processes X and X' defined for $K - N \leq n \leq K$, such that X evolves *via* the adaptive transition kernels P_{Γ_n}, while X' evolves *via* the fixed kernel $P' = P_{\Gamma_{K-N}}$. With probability at least $1 - 2\varepsilon$, the two processes remain equal for all times n with $K - N \leq n \leq K$. Hence, their un-coupling probability over this time interval is bounded above by 2ε. Consequently, conditional on X_{K-N} and Γ_{K-N}, the law of X_K lies within 2ε (measured in total variation distance) of the law of X'_K, which in turn lies within ε of the stationarity distribution π. Hence, the law of X_K is within 3ε of π. Since this holds for any $\varepsilon > 0$ (for sufficiently large $K = K(\varepsilon)$), it follows that the law of X_K converges to π as $K \to \infty$. Accordingly the adaptive process X is indeed a "valid" Monte Carlo algorithm for approximately sampling from π; namely it converges asymptotically to π. The proof of a more general result (Roberts and Rosenthal, 2007, Theorem 13), is quite similar, only requiring one additional ε.

3 *MEXIT* for discrete-time countable state-space

Having motivated the prominence of un-coupling arguments in key statistical and probabilistic settings, we now turn to an explicit construction of *MEXIT*. We begin by considering two discrete-time Markov chains defined on the same countable discrete state-space, begun at the same initial state s_0. We suppose that these chains are governed by transition probability kernels $p(a, b)$ and $q(a, b)$, respectively.

We extend the state-space by keeping track of the past trajectory of each chain (its "genealogy"). The state of one of these chains at time n will thus be a sequence or genealogy $\mathbf{s} \cdot a = (s_0, s_1, \ldots, s_n)$ of $n + 1$ states. Let $\mathbf{s} \cdot a$ denote the sequence or genealogy $\mathbf{s} = (s_0, s_1, \ldots, s_n, a)$ of $n + 2$ states, corresponding to the chain moving to state a at time $n + 1$. When discussing probabilities of paths, we will employ the abbreviations $p(\mathbf{s}) = p(s_0, s_1)p(s_1, s_2) \ldots p(s_{n-1}, s_n)$, $q(\mathbf{s}) = q(t_0, t_1)q(t_1, t_2) \ldots q(t_{n-1}, t_n)$. Note that $p(\mathbf{s} \cdot a) = p(\mathbf{s})p(s_n, a)$ *et cetera*. We define a coupling between the two processes as a random process on the cartesian product of the (extended) state-space with itself, whose marginal distributions are those of the individual processes.

Definition 2 (Coupling of two Markov chains)**.** A *coupling* of two (genealogical) Markov chains is a random process (*not* necessarily Markov!) with states $(\mathbf{s}, \mathbf{t})$ at time n given by pairs of genealogies of length n, such that if the probability of seeing state $(\mathbf{s}, \mathbf{t})$ at time n is given by $r(\mathbf{s}, \mathbf{t})$, then

$$\sum_{\mathbf{t}} r(\mathbf{s}, \mathbf{t}) \quad = \quad p(\mathbf{s}) \qquad \text{(row-marginals)}\,, \tag{1}$$

$$\sum_{\mathbf{s}} r(\mathbf{s}, \mathbf{t}) \quad = \quad q(\mathbf{t}) \qquad \text{(column-marginals)}\,. \tag{2}$$

Moreover, probabilities at consecutive times are related, for example, by

$$\sum_{a} \sum_{b} r(\mathbf{s} \cdot a, \mathbf{t} \cdot b) \quad = \quad r(\mathbf{s}, \mathbf{t}) \qquad \text{(inheritance)}\,. \tag{3}$$

Remark 3. A coupling of two non-genealogical Markov chains can be converted into the above form simply by keeping track of the genealogies.

Remark 4. Both chains begin at a fixed starting point s_0, so we set $p((s_0)) = q((s_0)) = 1$. The common starting point exactly means that the chains initially have the same trajectory. *MEXIT* occurs when first the trajectories split apart and disagree: the tree-like nature of genealogical state-space means the chains can never recombine.

Remark 5. Astute readers will already be asking whether, in the light of the genealogical nature of Definition 2 the processes really need to be Markov chains. In fact they do not, but we defer consideration of this point to Section 3, where we generalize the whole treatment to apply to rather general random processes both in discrete and in continuous time.

A *MEXIT* coupling is one which achieves the bound prescribed by the Aldous (1983) coupling inequality (Lemma 3.6 therein), thus (stochastically) maximising the time at which the chains split apart.

Definition 6 (*MEXIT* coupling)**.** Suppose that the *MEXIT equation* holds for all genealogical states $\mathbf{s}$:

$$r(\mathbf{s}, \mathbf{s}) \quad = \quad p(\mathbf{s}) \wedge q(\mathbf{s})\,. \tag{4}$$

Then the coupling is a *maximal exit coupling* (*MEXIT* coupling).

We now prove that *MEXIT* couplings always exist.

Theorem 7. *Consider two discrete-time Markov chains taking values in a given countable state-space and started at the same initial state s_0. A* MEXIT *coupling can be constructed so that the coupling moves as a Markov chain killed at* MEXIT*, and then the two separated components move independently as Markov chains conditioned so that the pair of Markov chains avoids certain "zero-regions".*

Proof. The pre-*MEXIT* component is easily described. The pre-*MEXIT* chain can be viewed as a single random process which is killed on *MEXIT*. In fact it can be viewed as a time-inhomogeneous genealogical Markov chain governed (up to killing time) by a time-inhomogeneous sub-Markovian kernel

$$\rho_{\text{pre-}MEXIT}(\mathbf{s}, \mathbf{s} \cdot a) \quad = \quad \frac{p(\mathbf{s} \cdot a) \wedge q(\mathbf{s} \cdot a)}{p(\mathbf{s}) \wedge q(\mathbf{s})}\,. \tag{5}$$

The unconditional *MEXIT* transition probabilities $p(\mathbf{s}) \wedge q(\mathbf{s})$, and $p(\mathbf{s}\cdot a) \wedge q(\mathbf{s}\cdot a)$ for states a, are related to the killing probability for this sub-Markovian kernel via:

$$p(\mathbf{s}) \wedge q(\mathbf{s}) - \sum_a p(\mathbf{s}\cdot a) \wedge q(\mathbf{s}\cdot a) = \left(1 - \sum_a \frac{p(\mathbf{s}\cdot a) \wedge q(\mathbf{s}\cdot a)}{p(\mathbf{s}) \wedge q(\mathbf{s})}\right) p(\mathbf{s}) \wedge q(\mathbf{s}) \quad (6)$$

$$= \left(1 - \sum_a \rho_{\text{pre-}MEXIT}(\mathbf{s}, \mathbf{s}\cdot a)\right) p(\mathbf{s}) \wedge q(\mathbf{s}). \quad (7)$$

Our task is to show one can specify non-Markovian couplings (given by $r(\mathbf{s}, \mathbf{t})$) such that $r(\mathbf{s}, \mathbf{s}) = p(\mathbf{s}) \wedge q(\mathbf{s})$ for all genealogical states $\mathbf{s}$. The issue is to show that we may construct the remaining probabilities $r(\mathbf{s}, \mathbf{t})$ $(\mathbf{s} \neq \mathbf{t})$ so as to have the correct row-marginals and column-marginals, while satisfying the inheritance property.

We proceed inductively. Certainly the assignation can be carried out for time 0, since both chains are then at state s_0. Consequently we can set $r((s_0), (s_0)) = 1$, and $r((s_0), (t_0)) = 0$ if $s_0 \neq t_0$, and thus row-marginals and column-marginals are all correct at time 0.

Proceeding with the inductive step, we suppose that at time n that we can assign $r(\mathbf{s}, \mathbf{t})$ for all paths $\mathbf{s}$ and $\mathbf{t}$ (up to time n) so as to obtain the correct row- and column-marginals and also so that the *MEXIT* equation (4) holds at time n. We lay out the $r(\mathbf{s}, \mathbf{t})$ in a matrix, with row-sums corresponding to the p-marginals, while the column-sums correspond to the q-marginals, as illustrated schematically in Figure 1. The induction step is successfully completed if we can assign values to the $r(\mathbf{s}\cdot a, \mathbf{s}\cdot b)$ satisfying the *MEXIT* equation (4) at time $n+1$ for $a = b$, such that the row-sums $\sum_b r(\mathbf{s}\cdot a, \mathbf{t}\cdot b)$ and the column-sums $\sum_a r(\mathbf{t}\cdot a, \mathbf{s}\cdot b)$ for $\mathbf{t} \neq \mathbf{s}$ are compatible with all the prescribed $r(\mathbf{s}, \mathbf{t})$. In that case, for example, one can use multiplication of the row-proportion (derived by analysis of the $(\mathbf{s}, \mathbf{s})$ cell on the diagonal) by that of the column-proportion (derived by analysis of the $(\mathbf{t}, \mathbf{t})$ cell on the diagonal) to determine the $r(\mathbf{s}\cdot a, \mathbf{t}\cdot b)$ for a specific $(\mathbf{s}, \mathbf{t})$ with $\mathbf{t} \neq \mathbf{s}$.

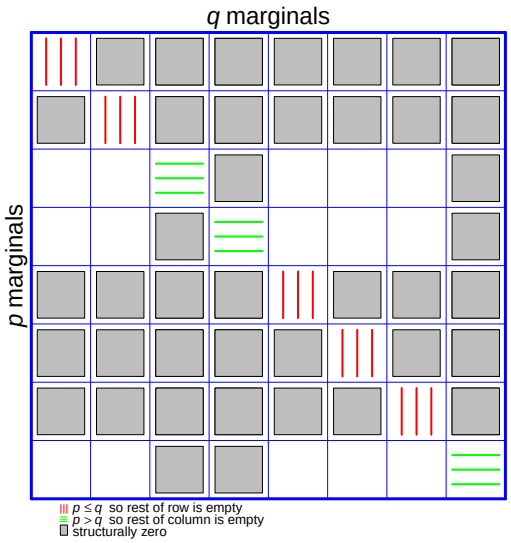

Figure 1: The $(\mathbf{s}, \mathbf{t})$ cells in a grid. For each cell $(\mathbf{s}, \mathbf{s})$ on the diagonal, if $p(\mathbf{s}) \leq q(\mathbf{s})$ or $p(\mathbf{s}) > q(\mathbf{t})^n$, respectively, the cells in the remainder of the row or the remainder of the column have zero probability.

Fix attention on a diagonal given by $(\mathbf{s}, \mathbf{s})$. Without loss of generality, suppose that $p(\mathbf{s}) \leq q(\mathbf{s})$, so that $r(\mathbf{s}, \mathbf{s}) = p(\mathbf{s})$. Note that this immediately implies that all other cells in this row (corresponding to $(\mathbf{s}, \mathbf{t})$ for $\mathbf{t} \neq \mathbf{s}$) must contain $r(\mathbf{s}, \mathbf{t}) = 0$. On the other hand, the residual column-sum for the other cells in the *column* (corresponding to $(\mathbf{t}, \mathbf{s})$ for $\mathbf{t} \neq \mathbf{s}$) must equal $q(\mathbf{s}) - p(\mathbf{s})$. For this cell, we separate the states corresponding to the next steps of the chains into two subsets:

$$A_+(\mathbf{s}) = \{a : p(\mathbf{s}\cdot a) \leq q(\mathbf{s}\cdot a)\},$$
$$A_-(\mathbf{s}) = \{a : p(\mathbf{s}\cdot a) > q(\mathbf{s}\cdot a)\}.$$

If $a \in A_+(\mathbf{s})$ then assign $r(\mathbf{s}\cdot a, \mathbf{s}\cdot a) = p(\mathbf{s}\cdot a)$, and observe that the $\mathbf{s}\cdot a$ sub-row must otherwise be filled with zeroes. Similarly, if $a \in A_-(\mathbf{s})$, then we assign $r(\mathbf{s}\cdot a, \mathbf{s}\cdot a) = q(\mathbf{s}\cdot a)$, and observe that the $\mathbf{s}\cdot a$ sub-*column* must otherwise be filled with zeroes. Accordingly the $A_+(\mathbf{s}) \times A_+(\mathbf{s})$ and $A_-(\mathbf{s}) \times A_-(\mathbf{s})$ sub-matrices are diagonal, while the $A_+(\mathbf{s}) \times A_-(\mathbf{s})$ sub-matrix is entirely filled with zeroes. The refined grid is illustrated schematically in Figure 2.

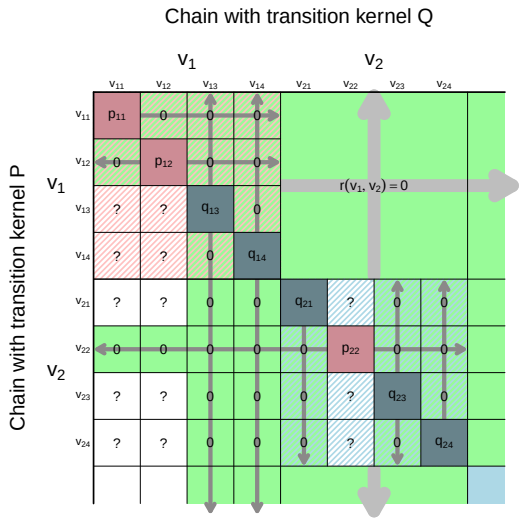

Figure 2: Refinement of the two neighboring cells into sub-cells $(\mathbf{s}\cdot a, \mathbf{s}\cdot b)$. Diagonal cells $(\mathbf{s}\cdot a, \mathbf{s}\cdot a)$ are assigned to $A_+(\mathbf{s})$ or to $A_-(\mathbf{s})$, according to whether $p(\mathbf{s}\cdot a) \le q(\mathbf{s}\cdot a)$ or $p(\mathbf{s}\cdot a) < q(\mathbf{s}\cdot a)$. Depending on which inequality holds, the sub-block cells in the remainder of the row or the remainder of the column hold zero probability.

Our task now reduces to showing that the $A_-(\mathbf{s}) \times A_+(\mathbf{s})$ sub-matrix and the column-sums $\sum_a r(\mathbf{t}\cdot a, \mathbf{s}\cdot b)$ for $\mathbf{t} \neq \mathbf{s}$ can be arranged to satisfy the row and column marginal requirements. We need not consider the row-sums $\sum_b r(\mathbf{s}\cdot a, \mathbf{t}\cdot b)$ for $\mathbf{t} \neq \mathbf{s}$; these must all vanish since $r(\mathbf{s}, \mathbf{s}) = p(\mathbf{s})$.

The row marginal requirements are satisfied if, for $a \in A_-(\mathbf{s})$, we can assign $\sum_{b\in A_+(\mathbf{s})} r(\mathbf{s}\cdot a, \mathbf{s}\cdot b)$ such that

$$p(\mathbf{s}\cdot a) \quad = \quad \sum_b r(\mathbf{s}\cdot a, \mathbf{s}\cdot b) \quad = \quad q(\mathbf{s}\cdot a) + \sum_{b\in A_+(\mathbf{s})} r(\mathbf{s}\cdot a, \mathbf{s}\cdot b)\,.$$

For $a \in A_-(\mathbf{s})$, we therefore require

$$\sum_{b\in A_+(\mathbf{s})} r(\mathbf{s}\cdot a, \mathbf{s}\cdot b) \quad = \quad p(\mathbf{s}\cdot a) - q(\mathbf{s}\cdot a)\,. \tag{8}$$

(Note: if $a \in A_-(\mathbf{s})$ then by definition $q(\mathbf{s}\cdot a) < p(\mathbf{s}\cdot a)$.) This will ensure that row-sum requirements are met for all the $\mathbf{s}\cdot a$ genealogical states. For $a \in A_+(\mathbf{s})$, this follows from the diagonal form of $A_+(\mathbf{s}) \times A_+(\mathbf{s})$, in addition to the fact that the $A_+(\mathbf{s}) \times A_-(\mathbf{s})$ sub-matrix is entirely filled with zeroes as well as the fact that $r(\mathbf{s}, \mathbf{t}) = 0$ if $\mathbf{t} \neq \mathbf{s}$, since $p(\mathbf{s}) \le q(\mathbf{s})$.

We complete the proof by showing that we can divide the row-sum in (8), and divide the $r(\mathbf{t}, \mathbf{s})$ (for $\mathbf{t} \neq \mathbf{s}$) among columns so as to ensure the column marginals are as required. We implement this as follows: column $\mathbf{s}\cdot b$ receives the proportion $\pi_1(\mathbf{s}\cdot b)$, where

$$\pi_1(\mathbf{s}\cdot b) \quad = \quad \begin{cases} \frac{q(\mathbf{s}\cdot b) - p(\mathbf{s}\cdot b)}{\sum_{c\in A_+(\mathbf{s})}(q(\mathbf{s}\cdot c) - p(\mathbf{s}\cdot c))} & \text{if } b \in A_+(\mathbf{s}) \\ 0 & \text{if } b \in A_-(\mathbf{s})\,. \end{cases} \tag{9}$$

To justify the above construction, observe that the total contribution from $A_-(\mathbf{s}) \times A_+(\mathbf{s})$ is $\sum_{a\in A_-(\mathbf{s})}(p(\mathbf{s}\cdot a) - q(\mathbf{s}\cdot a))$, while the total contribution from $(\mathbf{t}, \mathbf{s})$ for $\mathbf{t} \neq \mathbf{s}$) is $q(\mathbf{s}) - p(\mathbf{s})$. It follows that the $\mathbf{s}\cdot b$ column

obtains $\pi_1(\mathbf{s}\cdot b)$ of this sum together with a contribution $p(s\cdot b)$ from the diagonal matrix $A_+(\mathbf{s})\times A_+(\mathbf{s})$. This leads to a total $\mathbf{s}\cdot b$ column marginal of

$$\begin{aligned}
& p(s\cdot b)+\pi_1(\mathbf{s}\cdot b)\times\left(q(\mathbf{s})-p(\mathbf{s})+\sum_{a\in A_-(\mathbf{s})}(p(\mathbf{s}\cdot a)-q(\mathbf{s}\cdot a))\right)\\
= \;& p(s\cdot b)+\pi_1(\mathbf{s}\cdot b)\times\left(\sum_{a\in A_+(\mathbf{s})}(q(\mathbf{s}\cdot a)-p(\mathbf{s}\cdot a))\right)\\
= \;& p(s\cdot b)+\frac{q(\mathbf{s}\cdot b)-p(\mathbf{s}\cdot b)}{\sum_{c\in A_+(\mathbf{s})}(q(\mathbf{s}\cdot c)-p(\mathbf{s}\cdot c))}\times\left(\sum_{a\in A_+(\mathbf{s})}(q(\mathbf{s}\cdot a)-p(\mathbf{s}\cdot a))\right)\\
= \;& q(\mathbf{s}\cdot b),
\end{aligned}$$

as required.

Accordingly, we deduce that all marginal requirements are satisfied if we make the following assignations:

1. for $a\in A_-(\mathbf{s})$, $b\in A_+(\mathbf{s})$ assign

$$r(\mathbf{s}\cdot a,\mathbf{s}\cdot b) \;=\; \pi_1(\mathbf{s}\cdot b)\sum_{c\in A_-(\mathbf{s})}(p(\mathbf{s}\cdot c)-q(\mathbf{s}\cdot c))\,;$$

2. for $\mathbf{t}\neq\mathbf{s}$, if $r(\mathbf{s},\mathbf{s})=p(\mathbf{s})$ and $r(\mathbf{t},\mathbf{t})=q(\mathbf{t})$ (for otherwise the cell is assigned to be zero) assign

$$r(\mathbf{t}\cdot a,\mathbf{s}\cdot b) \;=\; \pi_1(\mathbf{s}\cdot b)\pi_2(\mathbf{s}\cdot a)r(\mathbf{t},\mathbf{s})\,.$$

Here the proportions $\pi_2(\mathbf{s}\cdot a)$ are defined by reversing roles of rows and columns, p and q:

$$\pi_2(\mathbf{s}\cdot b) \;=\; \begin{cases}\frac{p(\mathbf{s}\cdot b)-q(\mathbf{s}\cdot b)}{\sum_{c\in A_-(\mathbf{s})}(p(\mathbf{s}\cdot c)-q(\mathbf{s}\cdot c))} & \text{if } b\in A_-(\mathbf{s})\,.\\ 0 & \text{if } b\in A_+(\mathbf{s})\,.\end{cases} \tag{10}$$

This concludes the proof. □

Remark 8. We can (over-)parametrize all possible *MEXIT* couplings by replacing the assignations at (9) and (10) using copulae (Nelson, 2006) to parametrize the dependence between changes in the p-chain and the q-chain.

Having proven the existence of *MEXIT* couplings, we now provide calculations of *MEXIT* rate bounds (Subsection 3.1) and gain further insight into *MEXIT* by considering its connection with the Radon-Nikodym derivative (Subsection 3.2). We finish Section 3 on an applied note with a discussion of *MEXIT* times for MCMC algorithms (Subsection 3.3).

3.1 *MEXIT* rate bound

We now consider *MEXIT* rate bounds.

Proposition 9. *Consider the context of Theorem 7. Suppose we know that there is some $\delta>0$ such that either:*

(a) for all $\mathbf{s}$ *and* a,

$$\frac{p(\mathbf{s}\cdot a)/p(\mathbf{s})}{q(\mathbf{s}\cdot a)/q(\mathbf{s})} \;\geq\; 1-\delta$$

or

(b) for all $\mathbf{s}$ *and* a,

$$\frac{q(\mathbf{s}\cdot a)/q(\mathbf{s})}{p(\mathbf{s}\cdot a)/p(\mathbf{s})} \;\geq\; 1-\delta.$$

Then

$$\mathbb{P}[\text{MEXIT } \textit{at time } n+1 \mid \textit{no } \text{MEXIT } \textit{by time } n] \;\leq\; \delta.$$

Proof. Assume (a) (then (b) follows by symmetry). We obtain

$$\begin{aligned}
&\mathbb{P}[\text{no } \textit{MEXIT} \text{ by time } n+1 \,|\, \text{no } \textit{MEXIT} \text{ by time } n] \\
&= \frac{\sum_{\mathbf{s},a}[p(\mathbf{s}\cdot a) \wedge q(\mathbf{s}\cdot a)]}{\sum_{\mathbf{s}}[p(\mathbf{s}) \wedge q(\mathbf{s})]} \\
&\geq \frac{\sum_{\mathbf{s},a}[(1-\delta)q(\mathbf{s}\cdot a)\frac{p(\mathbf{s})}{q(\mathbf{s})} \wedge q(\mathbf{s}\cdot a)]}{\sum_{\mathbf{s}}[p(\mathbf{s}) \wedge q(\mathbf{s})]} \\
&= \frac{\sum_{\mathbf{s},a}\frac{q(\mathbf{s}\cdot a)}{q(\mathbf{s})}[(1-\delta)p(\mathbf{s}) \wedge q(\mathbf{s})]}{\sum_{\mathbf{s}}[p(\mathbf{s}) \wedge q(\mathbf{s})]} \\
&= \frac{\sum_{\mathbf{s}}[(1-\delta)p(\mathbf{s}) \wedge q(\mathbf{s})]}{\sum_{\mathbf{s}}[p(\mathbf{s}) \wedge q(\mathbf{s})]} \\
&\geq 1-\delta\,.
\end{aligned}$$

□

The above is the discrete state-space version of a bound contained in Völlering (2016). It should be noted that this bound applies equally well to faithful couplings, which typically degenerate in continuous time. (See Theorem 30 below for an example of this in the context of suitably regular diffusions.) Two corollaries of Proposition 9 follow immediately:

Corollary 10. *Under the conditions of Proposition 9,* $\mathbb{P}[\textit{no}\ \text{MEXIT}\ \textit{by time}\ n] \geq (1-\delta)^n$.

Corollary 11. *Under the conditions of Proposition 9,* $\mathbb{E}\,[\text{MEXIT}\ \textit{time}] \geq (1/\delta)$.

3.2 A Radon-Nikodym perspective on *MEXIT*

In this section, we explore a simple and natural connection of *MEXIT* to the value of the Radon-Nikodym derivative of q with respect to p.

In our discussion, it will suffice to consider *MEXIT* when the historical probability of the current path under both p and q are close to being equal, rare big jumps excepting. It follows from our *MEXIT* construction that the probability of *not* "MEXITing" by time n is equal to $\sum_{\mathbf{s}}(p(\mathbf{s}) \wedge q(\mathbf{s}))$, where the sum is over all length-n paths $\mathbf{s}$. Hence, conditional on having followed the path $\mathbf{s}$ up to time n and not "MEXITed," the conditional probability of *not* "MEXITing" at time $n+1$ is equal to

$$\frac{\sum_a(p(\mathbf{s}\cdot a) \wedge q(\mathbf{s}\cdot a))}{p(\mathbf{s}) \wedge q(\mathbf{s})}\,,$$

which follows from equations (6) and (7). Thus, the probability of "MEXITing" at time $n+1$ is

$$1 - \frac{\sum_a(p(\mathbf{s}\cdot a) \wedge q(\mathbf{s}\cdot a))}{p(\mathbf{s}) \wedge q(\mathbf{s})} = \frac{(p(\mathbf{s}) \wedge q(\mathbf{s})) - \sum_a(p(\mathbf{s}\cdot a) \wedge q(\mathbf{s}\cdot a))}{p(\mathbf{s}) \wedge q(\mathbf{s})}\,.$$

In particular, if $p(\mathbf{s}) > q(\mathbf{s})$ and $p(\mathbf{s}\cdot a) > q(\mathbf{s}\cdot a)$ for all a, then the numerator is zero, so the probability of "MEXITing" is zero. That is, "MEXITing" can only happen when the relative ordering of $(p(\mathbf{s}), q(\mathbf{s}))$ and $(p(\mathbf{s}\cdot a), q(\mathbf{s}\cdot a))$ are different.

We now rephrase the above arguments in the language of Radon-Nikodym derivatives. Let $q(a|\mathbf{s}) = q(\mathbf{s}\cdot a)/q(a)$, and $R(\mathbf{s}) = p(\mathbf{s})/q(\mathbf{s})$. Then the non-*MEXIT* probability is

$$\frac{\sum_a(p(\mathbf{s}\cdot a) \wedge q(\mathbf{s}\cdot a))}{p(\mathbf{s}) \wedge q(\mathbf{s})} = \mathbb{E}_{q(a|\mathbf{s})}\left[\frac{R(\mathbf{s}\cdot a) \wedge 1}{R(\mathbf{s}) \wedge 1}\right] = \mathbb{E}_{p(a|\mathbf{s})}\left[\frac{R(\mathbf{s}\cdot a)^{-1} \wedge 1}{R(\mathbf{s})^{-1} \wedge 1}\right].$$

Note that $\mathbb{E}_{q(a|\mathbf{s})}\left[R(\mathbf{s}\cdot a)\right] = R(\mathbf{s})$. Thus, if we have either $R(\mathbf{s}) < 1$ and $R(\mathbf{s}\cdot a) < 1$ for all a, or $R(\mathbf{s}) > 1$ and $R(\mathbf{s}\cdot a) > 1$ for all a, then this non-*MEXIT* probability is one and thus the *MEXIT* probability is zero. That is, *MEXIT* can only occur when the Radon-Nikodym derivative R changes from more than 1 to less than 1 or vice-versa.

3.2.1 An example: *MEXIT* for simple random walks

To further elucidate the connection of *MEXIT* with the Radon-Nikodym derivative, we consider a concrete example: two simple random walks. Let "p" be simple random walk with up probability $\eta < 1/2$ and down probability $1-\eta$. Similarly, let "q" be a simple random walk with up probability $1-\eta$ and down probability η. *MEXIT* only occurs at 0, which is the only point at which the two processes have the same transition probability, i.e. at which the Radon-Nikodym derivative is equal to 1. Indeed, the "pre-*MEXIT*" process (i.e., the joint process, conditional on *MEXIT* not having yet occurred) evolves with the following dynamics:

- For $k > 0$, $P(k, k+1) = \eta$, and $P(k, k-1) = 1-\eta$.
- For $k < 0$, $P(k, k+1) = 1-\eta$, and $P(k, k-1) = \eta$.
- $P(0,1) = P(0,-1) = \eta$ with *MEXIT* probability $1-2\eta$ when we are at 0.

Note that the chain P is defective at 0, but otherwise has a drift towards the *MEXIT* point 0. Letting Q_t denote the number of times the joint process (conditional on not yet "MEXITing") hits 0 up to and including time t, we have that

$$\mathbb{P}[\textit{MEXIT} \text{ by time } t \mid Q_{t-1}] \quad = \quad 1 - (2\eta)^{Q_{t-1}}. \tag{11}$$

3.3 An application: noisy MCMC

The purpose of this section is to provide an application of *MEXIT* for discrete-time countable state-spaces. We do so by comparing the *MEXIT* time τ of the *penalty method* MCMC algorithm with the usual Metropolis-Hastings algorithm.

In the usual Metropolis-Hastings algorithm, starting at a state X, we propose a new state Y, and then accept it with probability $1 \wedge A(X,Y)$, where $A(X,Y)$ is an appropriate acceptance probability formula. In *noisy MCMC* (specifically, the *penalty method* MCMC, see Ceperley and Dewing (1999); Nicholls, Fox, and Watt (2012); Medina-Aguayo, Lee, and Roberts (2015); Alquier, Friel, Everitt, and Boland (2016)) which is similar to but different from the *pseudo-marginal MCMC* method of Andrieu and Roberts (2009)), we accept with probability $\widehat{\alpha}(X,Y) := 1 \wedge (A(X,Y)\, W)$, where W is an independent random variable. For simplicity, assume that $W = \exp(N)$ where $N \sim \text{Normal}(-\sigma^2/2, \sigma^2)$ for some fixed $\sigma > 0$ (so that $\mathbb{E}[W] = \mathbb{E}[\exp(N)] = 1$), i.e. that $\widehat{\alpha}(X,Y) := 1 \wedge (A(X,Y)\exp(N))$. We now show that the *penalty method MCMC* produces a Metropolis-Hastings algorithm with sub-optimal acceptance probability.

Proposition 12. *The penalty method MCMC produces a Metropolis-Hastings algorithm with (sub-optimal) acceptance probability* $\widetilde{\alpha}(X,Y,\sigma) := \mathbb{E}[\widehat{\alpha}(X,Y) \mid X, Y]$ *given by*

$$\widetilde{\alpha}(X,Y,\sigma) \quad = \quad \Phi\left[\frac{\log A(X,Y)}{\sigma} - \frac{\sigma}{2}\right] + A(X,Y)\,\Phi\left[-\frac{\sigma}{2} - \frac{\log A(X,Y)}{\sigma}\right].$$

Proof. We invoke Proposition 2.4 of Roberts, Gelman, and Gilks (1997), which states that if $B \sim \text{Normal}(\mu, \sigma^2)$, then

$$\mathbb{E}\left[1 \wedge e^B\right] \quad = \quad \Phi\left(\frac{\mu}{\sigma}\right) + \exp(\mu + \sigma^2/2)\,\Phi\left[-\sigma - \frac{\mu}{\sigma}\right].$$

Note

$$\widetilde{\alpha}(X,Y,\sigma) \quad = \quad \mathbb{E}[\widehat{\alpha}(X,Y)] \quad = \quad \mathbb{E}\left[1 \wedge (A(X,Y)e^N)\right] \quad = \quad \mathbb{E}\left[1 \wedge e^{N(-\sigma^2/2 + \log A(X,Y),\ \sigma^2)}\right].$$

After straightforward algebra, the right-hand side of the last equality simplifies to

$$\Phi\left[\frac{\log A(X,Y)}{\sigma} - \frac{\sigma}{2}\right] + A(X,Y)\,\Phi\left[-\frac{\sigma}{2} - \frac{\log A(X,Y)}{\sigma}\right].$$

□

Proposition 13. $A(X,Y)\,\phi\left[-\frac{\sigma}{2} - \frac{\log A(X,Y)}{\sigma}\right] \quad = \quad \phi\left[\frac{\log A(X,Y)}{\sigma} - \frac{\sigma}{2}\right]$.

Proof. We calculate

$$\begin{aligned} & A(X,Y)\,\phi\left[-\frac{\sigma}{2}-\frac{\log A(X,Y)}{\sigma}\right] \\ = \quad & \frac{1}{\sqrt{2\pi}}\exp(\log A(X,Y)-\frac{1}{2}\left(-\frac{\sigma}{2}-\left(\frac{\log A(X,Y)}{\sigma}\right)^2\right) \\ = \quad & \frac{1}{\sqrt{2\pi}}\exp\left(-\frac{1}{2}\left(\frac{\log A(X,Y)}{\sigma}-\frac{\sigma}{2}\right)^2\right) \\ = \quad & \phi\left(\frac{\log A(X,Y)}{\sigma}-\frac{\sigma}{2}\right). \end{aligned}$$

□

Proposition 14. *For any $a, s > 0$, we have that*

$$\frac{1}{a}\,\phi\left(\frac{\log a}{s}-\frac{s}{2}\right) \quad \leq \quad \frac{1}{\sqrt{2\pi}}. \tag{12}$$

Proof. This follows from noting

$$\begin{aligned} & \frac{1}{a}\,\phi\left(\frac{\log a}{s}-\frac{s}{2}\right) \\ = \quad & \frac{1}{\sqrt{2\pi}}\exp\left(-\log a-\frac{1}{2}\left(\frac{\log a}{s}-\frac{s}{2}\right)^2\right) \\ = \quad & \frac{1}{\sqrt{2\pi}}\exp\left(-\frac{1}{2}\left(\frac{\log a}{s}+\frac{s}{2}\right)^2\right) \leq \frac{1}{\sqrt{2\pi}}. \end{aligned}$$

□

Let $r(X)$ and $\widetilde{r}(X)$ be the probabilities of rejecting the proposal when starting at X for the original Metropolis-Hastings algorithm and the *penalty method* MCMC, respectively. We now proceed with Proposition 15.

Proposition 15. *For all X, Y in the state space, and $\sigma \geq 0$, the following seven statements hold*

(1) $\widetilde{\alpha}(X,Y) \leq \alpha(X,Y)$.
(2) $\widetilde{r}(X) \geq r(X)$.
(3) $\lim_{\sigma \searrow 0} \widetilde{\alpha}(X,Y,\sigma) = \alpha(X,Y)$.
(4) $\frac{d}{d\sigma}\widetilde{\alpha}(X,Y,\sigma) = -\phi\left[\frac{\log A(X,Y)}{\sigma}-\frac{\sigma}{2}\right]$.
(5) $0 \geq \frac{d}{d\sigma}\widetilde{\alpha}(X,Y,\sigma) \geq -1/\sqrt{2\pi}$.
(6) $\widetilde{\alpha}(X,Y,\sigma) \geq \alpha(X,Y) - \sigma/\sqrt{2\pi}$.
(7) $\frac{\widetilde{\alpha}(X,Y,\sigma)}{\alpha(X,Y)} \geq 1 - \sigma/\sqrt{2\pi}$.

Proof. For statement (1), apply Jensen's inequality. Note that

$$\begin{aligned} & \mathbb{E}\left[\widetilde{\alpha}(X,Y)\,|\,X,Y\right] \\ = \; & \mathbb{E}\left[1 \wedge (A(X,Y)e^N)\,|\,X,Y\right] \\ \leq \; & 1 \wedge \mathbb{E}\left[(A(X,Y)e^N)\right] \\ = \; & 1 \wedge (A(X,Y)\,\mathbb{E}\left[e^N\right]) \\ = \; & 1 \wedge A(X,Y) = \alpha(X,Y)\,. \end{aligned}$$

Statement (2) follows immediately from statement (1) by taking the complements of the expectations of the $\alpha(X,Y)$ and $\widetilde{\alpha}(X,Y)$ with respect to Y.

For statement (3), note that if $A(X,Y) > 1$ then $\lim_{\sigma \searrow 0} \widetilde{\alpha}(X,Y,\sigma) = \Phi[+\infty] + A(X,Y)\,\Phi[-\infty] = 1$, while if $A(X,Y) < 1$ then $\lim_{\sigma \searrow 0} \widetilde{\alpha}(X,Y,\sigma) = \Phi[-\infty] + A(X,Y)\,\Phi[+\infty] = 0 + A(X,Y)\,1 = A(X,Y)$. Further,

if $A(X,Y) = 1$ then $\lim_{\sigma \searrow 0} \widetilde{\alpha}(X,Y,\sigma) = \Phi[0] + A(X,Y)\,\Phi[0] = (1/2) + (1)(1/2) = 1$. Thus, in all cases, $\lim_{\sigma \searrow 0} \widetilde{\alpha}(X,Y,\sigma) = 1 \wedge A(X,Y) = \alpha(X,Y)$.

For statement (4), we use Proposition 13 to compute

$$\begin{aligned}
& \frac{d}{d\sigma}\widetilde{\alpha}(X,Y,\sigma) \\
= \;& \frac{d}{d\sigma}\left(\Phi\left[\frac{\log A(X,Y)}{\sigma} - \frac{\sigma}{2}\right] + A(X,Y)\,\Phi\left[-\frac{\sigma}{2} - \frac{\log A(X,Y)}{\sigma}\right]\right) \\
= \;& \phi\left[\frac{\log A(X,Y)}{\sigma} - \frac{\sigma}{2}\right]\left(-\frac{\log A(X,Y)}{\sigma^2} - \frac{1}{2}\right) + A(X,Y)\;\phi\left[-\frac{\sigma}{2} - \frac{\log A(X,Y)}{\sigma}\right] \\
= \;& -\frac{1}{2} + \frac{\log A(X,Y)}{\sigma^2} \;=\; -\phi\left[\frac{\log A(X,Y)}{\sigma} - \frac{\sigma}{2}\right].
\end{aligned}$$

Since $0 \le \phi(\cdot) \le \frac{1}{\sqrt{2\pi}}$, statement (5) follows immediately. Statement (6) then follows by integrating from 0 to σ. For statement (7), note that if $A(X,Y) \ge 1$ then $\alpha(X,Y) = 1$ and the result then follows from statement (6). If instead $A(X,Y) < 1$, then $\alpha(X,Y) = A(X,Y)$, and we may invoke Proposition 14 to obtain

$$\begin{aligned}
& \frac{\widetilde{\alpha}(X,Y,\sigma)}{\alpha(X,Y)} \\
= \;& 1 - \frac{\alpha(X,Y) - \widetilde{\alpha}(X,Y,\sigma)}{\alpha(X,Y)} \\
= \;& 1 - \int_{u=0}^{\sigma} \frac{1}{\alpha(X,Y)}\left(\frac{d}{du}\widetilde{\alpha}(X,Y,u)\right)du \\
= \;& 1 - \int_{u=0}^{\sigma} \frac{1}{A(X,Y)}\phi\left[\frac{\log A(X,Y)}{\sigma} - \frac{\sigma}{2}\right]du \\
\ge \;& 1 - \int_{u=0}^{\sigma} \frac{1}{\sqrt{2\pi}}\,du \;=\; 1 - \frac{\sigma}{\sqrt{2\pi}}.
\end{aligned}$$

This concludes the proof. □

Let P be the law of a Metropolis-Hastings algorithm, and $\widetilde{P}$ the law of a corresponding noisy MCMC. We now prove Proposition 16 below, whose Corollary 17 uses *MEXIT* to control the discrepancy between the Metropolis-Hastings algorithm and the noisy MCMC algorithm.

Proposition 16.

$$\frac{d\widetilde{P}^{t+1}(\mathbf{s}\cdot a)}{dP^{t+1}(\mathbf{s}\cdot a)} \;\ge\; \frac{d\widetilde{P}^{t}(\mathbf{s})}{dP^{t}(\mathbf{s})}\left(1 - \frac{\sigma}{\sqrt{2\pi}}\right).$$

Proof. Note first that $\frac{d\widetilde{P}^t(\mathbf{s})}{dP^t(\mathbf{s})} = \gamma_1\gamma_2\ldots\gamma_n$ where each γ_i equals either $\frac{\widetilde{\alpha}(X_{i-1},X_i)}{\alpha(X_{i-1},X_i)}$ if the move from X_{i-1} to X_i is accepted and otherwise $\frac{\widetilde{r}(X)}{r(X)}$ if the move is rejected. Statement (2) of Proposition 15 tells us that, if we reject,

$$\frac{d\widetilde{P}^{t+1}(\mathbf{s}\cdot a)}{dP^{t+1}(\mathbf{s}\cdot a)} \;\ge\; \frac{d\widetilde{P}^{t}(\mathbf{s})}{dP^{t}(\mathbf{s})} \;\ge\; \frac{d\widetilde{P}^{t}(\mathbf{s})}{dP^{t}(\mathbf{s})}\left(1 - \frac{\sigma}{\sqrt{2\pi}}\right).$$

However, if we accept, then by statement (7) in Proposition 15, $\frac{d\widetilde{P}^{t+1}(\mathbf{s}\cdot a)}{dP^{t+1}(\mathbf{s}\cdot a)} \ge \frac{d\widetilde{P}^{t}(\mathbf{s})}{dP^{t}(\mathbf{s})}(1 - \frac{\sigma}{\sqrt{2\pi}})$, as claimed. □

The following Corollary to Proposition 16 now follows immediately.

Corollary 17. $\frac{d\widetilde{P}^t(\mathbf{s})}{dP^t(\mathbf{s})} \;\ge\; \left(1 - \frac{\sigma}{\sqrt{2\pi}}\right)^t$.

Applying Proposition 16 to Proposition 9 in Subsection 3.1, with $\delta = \frac{\sigma}{\sqrt{2\pi}}$, the following Corollary follows immediately.

Corollary 18. *The* MEXIT *time τ of the above penalty method MCMC algorithm, compared to the regular Metropolis-Hastings algorithm, satisfies the following two inequalities:*

$$\mathbb{P}[\tau > n] \quad \geq \quad \left(1 - \frac{\sigma}{\sqrt{2\pi}}\right)^n$$

and

$$\mathbb{E}[\tau] \quad \geq \quad \sqrt{2\pi}/\sigma.$$

4 *MEXIT* for general random processes

The methods and results of Section 3 can be generalized to the case when the two processes are general time-inhomogeneous random processes in discrete time with countable state-space. We now show that the same approach works for the case of general measurable state space $(E, \mathcal{E})$. The chief technical difficulty lies in expressing the calculations of Section 3 in measure-theoretic form. Note that in general the diagonal set $\Delta = \{(x, x) : x \in E\} \subset E \times E$ does not belong to the product σ-algebra $\mathcal{E} * \mathcal{E}$ (for a few examples of this phenomenon, see Stoyanov, 1997, Subsection 1.6). However Δ *is* analytic and thus universally measurable. Consequently, once we have determined a coupling probability measure on the product measure space $(E \times E, \mathcal{E} * \mathcal{E})$, then the diagonal set Δ will belong to the completion of the product σ-algebra $\mathcal{E} * \mathcal{E}$ under the coupling probability measure. More information about this aspect of measure-theoretic probability can be found, for example, in Dellacherie and Meyer (1979). We evade many measure-theoretic issues by working directly with probability measures.

4.1 Case of one time-step

To establish notation, we first review the simplest case of just one time-step. Consider two E-valued random variables X_1^+ and X_1^-, measurable with respect to $\mathcal{E}$ on E, with distributions $\mathcal{L}\left(X_1^+\right) = \mu_1^+$ and $\mathcal{L}\left(X_1^-\right) = \mu_1^-$ on $(E, \mathcal{E})$. We recall that the *meet measure* $\hat{\mu}_1 = \mu_1^+ \wedge \mu_1^-$ in the lattice of non-negative measures on $(E, \mathcal{E}_1)$ can be described explicitly using the Hahn-Jordan decomposition (Halmos, 1978, §28) as

$$\mu_1^+ - \mu_1^- \quad = \quad \nu_1^+ - \nu_1^- \tag{13}$$

for unique non-negative measures ν_1^+ and ν_1^- of disjoint support. The condition of disjoint support implies that

$$\hat{\mu}_1 \quad = \quad \mu_1^+ - \nu_1^+ \quad = \quad \mu_1^- - \nu_1^- \tag{14}$$

is the maximal non-negative measure $\widetilde{\mu}$ such that

$$\widetilde{\mu}(D) \quad \leq \quad \min\{\mu_1^+(D), \mu_1^-(D)\} \qquad \text{for all } D \in \mathcal{E}\,.$$

More explicitly, note that the probability measures $\mu_1^\pm$ are both absolutely continuous with respect to their average, $\overline{\mu}_1 = \frac{1}{2}(\mu_1^+ + \mu_1^-)$. This permits employment of the Radon-Nikodym theorem: the equations

$$f_1^\pm \quad = \quad \mathrm{d}\,\mu_1^\pm / \,\mathrm{d}\,\overline{\mu}_1 \tag{15}$$

determine non-negative densities $f_1^\pm$ for $\mu_1^\pm$. The densities are defined up to sets of $\overline{\mu}_1$-measure zero, and we suppose specific choices of non-negative representatives are made within the respective equivalence classes. The disjoint supports of the $\nu_1^\pm$ can then be defined by $[f_1^\pm > 1]$. Note also that $f_1^+ + f_1^- = 2$, and that $\mathrm{d}\,\hat{\mu}_1 = (f_1^+ \wedge f_1^-)\,\mathrm{d}\,\overline{\mu}_1$, and that

$$\hat{\mu}_1(D) \quad = \quad \mu_1^\pm(D \cap [f_1^+ < 1]) + \mu_2^\pm(D \cap [f_1^+ \geq 1]) \qquad \text{for all } D \in \mathcal{E}\,.$$

It is useful to note that

$$\mathrm{d}\,\nu_1^\pm \quad = \quad ((f_1^\pm - f_1^\mp) \vee 0)\;\; \mathrm{d}\,\overline{\mu}_1\,. \tag{16}$$

Lemma 19. *The simplest* MEXIT *problem, concerning two random variables X_1^+ and X_1^- taking values in the same measurable space $(E, \mathcal{E})$, is solved by maximal coupling of the two marginal probability measures $\mu_1^+ = \mathcal{L}\left(X_1^+\right)$ and $\mu_1^- = \mathcal{L}\left(X_1^-\right)$ using a joint probability measure m_1 on the product measure space $(E \times E, \mathcal{E} * \mathcal{E})$ such that*

1. *m_1 has marginal distributions μ_1^+ and μ_1^- on the two coordinates,*

2. *$m_1 \geq \imath_{\Delta *}\hat{\mu}_1$, where the non-negative measure $\hat{\mu}_1 = \mu_1^+ \wedge \mu_1^-$ is the meet measure for μ_1^+ and μ_1^-, and $\imath_{\Delta *}$ is the push-forward map corresponding to the $(\mathcal{E} : \mathcal{E} * \mathcal{E})$-measurable "diagonal injection" $\imath_\Delta : E \to E \times E$ given by $\imath_\Delta(x) = (x, x)$.*

Proof. One possible explicit construction for m_1 is

$$m_1 \quad = \quad \imath_{\Delta *}\hat{\mu}_1 + \frac{1}{\nu_1^+(E)}\nu_1^+ \otimes \nu_1^- \,, \tag{17}$$

where $\nu_1^+ \otimes \nu_1^-$ is the product measure on $(E \times E, \mathcal{E} * \mathcal{E})$. (Remark 20 discusses other possibilities). It follows directly from (13) that $\nu_1^+(E) = \nu_1^-(E)$. Maximality of the coupling (which is to say, maximality of $m_1(\Delta) = \hat{\mu}_1(E)$ compared to all other probability measures with these marginals) follows from maximality of the meet measure. This completes the proof. □

Remark 20. The maximal coupling (17) is by no means unique: the scaled product measure in (17) can be replaced by any other bivariate measure Γ_1 in the *Fréchet class* $F(\nu_1^+, \nu_1^-)$ of non-negative measures on $(E \times E, \mathcal{E} * \mathcal{E})$ with marginals ν_1^+ and ν_1^-; and the general maximal coupling can then be expressed as

$$m_1 \quad = \quad \imath_{\Delta *}\hat{\mu}_1 + \Gamma_1 \tag{18}$$

for specified $\Gamma_1 \in F(\nu_1^+, \nu_1^-)$. (Note that $F(\nu_1^+, \nu_1^-)$ is non-void only when, as in this case, the total masses $\nu_1^+(E) = \nu_1^-(E)$ agree.) Moreover the Fréchet class $F(\nu_1^+, \nu_1^-)$ can be (over-)parametrized by the corresponding class of copulae (Nelson, 2006) using inverse distribution transform techniques.

Given this construction, we can realize X_1^+ and X_1^- as the coordinate maps for $E \times E$. The diagonal set Δ lies in the σ-algebra obtained by completing $\mathcal{E} * \mathcal{E}$ under m_1. Consequently the probability statements

$$\mathbb{P}\left[X_1^+ \in D \,;\, X_1^+ = X_1^-\right] \quad = \quad \hat{\mu}_1(D) \qquad \text{for all } D \in \mathcal{E} \tag{19}$$

hold for any maximal coupling of X_1^+ and X_1^-.

It is convenient at this point to note a quick way to recognize when a given coupling is maximal.

Lemma 21 (Recognition Lemma for Maximal Coupling)**.** *Given a coupling probability measure m^* for $\mathcal{E}$-measurable random variables X_1^+ and X_1^- (with distributions $\mathcal{L}\left(X_1^+\right) = \mu_1^+$ and $\mathcal{L}\left(X_1^-\right) = \mu_1^-$), this coupling is maximal if the two non-negative measures*

$$\nu_1^{\pm,*} \,:\, D \quad \mapsto \quad m^*[X_1^\pm \in D \,;\, X_1^+ \neq X_1^-] \tag{20}$$

(defined for $D \in \mathcal{E}$) have disjoint supports. Moreover in this case the meet measure for the two probability distributions $\mathcal{L}\left(X_1^+\right)$ and $\mathcal{L}\left(X_1^-\right)$ is given by

$$\hat{\mu}_1(D) \quad = \quad m^*[X_1^+ \in D \,;\, X_1^+ = X_1^-] \qquad \textit{for all } D \in \mathcal{E}\,. \tag{21}$$

Proof. This follows immediately from the uniqueness of the non-negative measures $\nu_1^\pm$ of disjoint support appearing in the Hahn-Jordan decomposition, since a sample-wise cancellation of events shows that

$$\mu_1^+ - \mu_1^- \quad = \quad \mathcal{L}\left(X_1^+\right) - \mathcal{L}\left(X_1^-\right) \quad = \quad \nu_1^{+,*} - \nu_1^{-,*}\,.$$

□

4.2 Case of n time-steps

We now state and prove the generalization of Theorem 7 to the case of discrete-time random processes taking values in a general state-space. The proof strategy follows that of Theorem 7; however, care needs to be taken to express concepts accurately in this general framework.

Theorem 22. *Consider two discrete-time random processes X^+ and X^-, begun at the same fixed initial point, taking values in a measurable state-space $(E, \mathcal{E})$ and run up to a finite time n. Maximal* MEXIT *couplings exist.*

Remark 23. One might suppose one needs $\mathcal{E} * \mathcal{E}$ measurability of the diagonal, so that the event of coupling is measurable. However the diagonal set is analytic and thus universally measurable. In particular, it is measurable with respect to σ-algebra completion of $\mathcal{E} * \mathcal{E}$ with respect to the coupling measure. In fact, the measurable state-space can be entirely general; in particular we do not require it to be Polish (in contrast to Völlering, 2016). The gain in generality may not be practically significant; however, it does focus attention on a constructive proof strategy.

Proof. Denote the two random processes by X^+ and X^-. Consider their trajectories from time 0 until time n. Since the trajectories start from the same fixed point in E, they can be regarded as random elements of the n-fold product of the original state-space. We denote this trajectory state-space (an n-fold Cartesian product of the original state-space) by E_n, with product σ-algebra $\mathcal{E}_n$.

We seek a probability measure m_n on the product measure space $(E_n \times E_n, \mathcal{E}_n * \mathcal{E}_n)$ such that its marginals have the distributions μ_n^+ and μ_n^- of the sequence formed by the first n steps of X^+ and X^-, respectively, and such that the m_n-measure of the event that *MEXIT* occurs after time k is maximal for each of the times $k = 1, \ldots, n$.

We have not required the processes to be stationary in time; accordingly it suffices to establish this result for $n = 2$, given a general maximal coupling of $X_\perp^+$ and $X_\perp^-$, after which a simple induction argument will yield the general result. Note that the natural projection $\Pi : E_2 \to E_1$ allows us to abbreviate notation by supposing that $\mathcal{E}_1 \leq \mathcal{E}_2$ as σ-algebras, and also by viewing functions on E_1 (such as $f_1^\pm$) as equally defined on E_2 *via* the corresponding projections (effectively, we abbreviate $f_1^\pm \circ \Pi$ by $f_1^\pm$).

Together with $\mu_1^\pm$, $\hat{\mu}_1$, $\overline{\mu}_1$, $\nu_1^\pm$, $f_1^\pm$ defined as above, set $\mu_2^\pm = \mathcal{L}\left((X_1^\pm, X_n^\pm)\right)$, and define $\hat{\mu}_2$, $\overline{\mu}_2$, $\nu_2^\pm$, $f_2^\pm$ correspondingly, defined for the measurable space $(E_2, \mathcal{E}_2)$.

The given maximal coupling on $E_1 \times E_1$ is denoted in the manner of (18) by

$$m_1 \quad = \quad \imath_{\Delta_1 *} \hat{\mu}_1 + \Gamma_1 \,,$$

where Γ_1 lies in the Fréchet class $F(\nu_1^+, \nu_1^-)$, and $\imath_{\Delta_1} : E_1 \to E_1 \times E_1$ is the $\mathcal{E}_1 : \mathcal{E}_1 * \mathcal{E}_1$-measurable diagonal map. Let $\Pi_1^+, \Pi_1^- : E_1 \times E_1 \to E_1$ be the projections onto first and second coordinates of $E_1 \times E_1$: then it is immediate from (14) that $\Pi_1^\pm{}_* m_1 = \mu_1^\pm$. We seek a coupling probability measure m_2 on $E_2 \times E_2$ which satisfies the marginal constraints $\Pi_2^\pm{}_* m_2 = \mu_2^\pm$ (for $\Pi_2^\pm : E_2 \times E_2 \to E_2$ the projections onto first and second coordinates of $E_2 \times E_2$) but also the inheritance constraint

$$m_2|_{\mathcal{E}_1 * \mathcal{E}_1} \quad = \quad m_1 \,.$$

We will need a probability kernel $p_{1\to 2}$ to link $\overline{\mu}_2 = \frac{1}{2}(\mu_2^+ + \mu_2^-)$ to $\overline{\mu}_1 = \frac{1}{2}(\mu_1^+ + \mu_1^-)$. Because $\overline{\mu}_2$ restricts on $\mathcal{E}_1$ to $\overline{\mu}_1$, we have the following absolute continuity relationship: for each fixed $D \in \mathcal{E}$, there is a $\mathcal{E}_1$-measurable density $p_{1\to 2}(x; D)$ giving

$$\overline{\mu}_2(D_1 \times D) \quad = \quad \int_{D_1} p_{1\to 2}(x; D)\, \overline{\mu}_1(\,\mathrm{d}\, x)$$

for all $D_1 \in \mathcal{E}_2$. Since $\overline{\mu}_2(D_1 \times D) \leq \overline{\mu}_1(D_1)$ for all $D_1 \in \mathcal{E}_1$, we can impose $0 \leq p_{1\to 2}(x; D) \leq 1$. Note in particular that we can take $p_{1\to 2}(x; E) = 1$. We have avoided the imposition of conditions such as Polish space properties which would guarantee existence of regular conditional probabilities, so we cannot deduce full σ-additivity in D for $p_{1\to 2}(x; D)$; however σ-additivity holds in $L^1(\overline{\mu}_1)$ and allows us to assert that

$$\int h(x_1, x_2)\, \overline{\mu}_2(\,\mathrm{d}(x_1, x_2)) \quad = \quad \int\int h(x_1, x_2) p_{1\to 2}(x_1; \,\mathrm{d}\, x_2)\, \overline{\mu}_1(\,\mathrm{d}\, x_1) \quad = \quad \int\int h \,\mathrm{d}(p_{1\to 2} * \overline{\mu}_1)\,, \tag{22}$$

when h is a bounded $\mathcal{E}_2$-measurable function and $x = (x_1, x_2) \in E_2$ where x_1 and x_2 correspond to the two components of $E_2 = E_1 \times E$. This is sufficient for our purposes, as our coupling constructions will always use the kernel to generate probability measures using integration. Note in particular that

$$\int f_2^\pm(x_1, x_2) p_{1\to 2}(x_1, \,\mathrm{d}\, x_2) \quad = \quad f_1^\pm(x_1) \tag{23}$$

follows from $\mu_2^\pm(D) = \mu_1^\pm(D)$ for all $D \in \mathcal{E}_1$.

We begin by noting that in general the Hahn-Jordan decomposition $\mu_2^+ - \mu_2^- = \nu_2^+ - \nu_2^-$ does not result in $\nu_2^\pm$ restricting to $\nu_1^\pm$ on $\mathcal{E}_1$. Indeed, from maximality of the meet measures,

$$\hat{\mu}_2(D) \quad \leq \quad \hat{\mu}_1(D) \qquad \text{for all } D \in \mathcal{E}_1\,;$$

moreover, by (14) and its analogue for $\hat{\mu}_2$, we may write the following relationship between non-negative measures determined on $\mathcal{E}_1$:

$$\eta_1 \quad = \quad \hat{\mu}_1 - \hat{\mu}_2|_{\mathcal{E}_1} \quad = \quad \nu_2^{\pm}|_{\mathcal{E}_1} - \nu_1^{\pm}\,. \tag{24}$$

The measure η_1 measures the amount of un-coupling that must happen between times 1 and 2. Note in particular that $\mathrm{d}\,\hat{\mu}_n = (f_n^{\pm} \wedge f_n^{\mp})\,\mathrm{d}\,\overline{\mu}_n$ and so we can use the probability kernel $p_{1\to 2}$ to deduce that

$$\mathrm{d}\,\eta_1 \quad = \quad \left((f_1^{\pm} \wedge f_1^{\mp}) - \int (f_2^{\pm} \wedge f_2^{\mp})\,\mathrm{d}\,p_{1\to 2}\right)\mathrm{d}\,\overline{\mu}_1\,. \tag{25}$$

These considerations allow us to write down an expression for a general maximal coupling measure m_2 at time 2 which is has the right marginals but also satisfies the m_1-inheritance requirement. We need

$$\begin{aligned} m_2 \quad &= \quad \mathrm{1}_{\Delta_2 *}\hat{\mu}_2 + m_2' + m_2''\,, \\ \Pi_2^{\pm}{}_* m_2 \quad &= \quad \mu_2^{\pm}\,, \\ m_2'|_{\mathcal{E}_1 * \mathcal{E}_1} \quad &= \quad \mathrm{1}_{\Delta_2 *}\eta_1\,, \\ m_2''|_{\mathcal{E}_1 * \mathcal{E}_1} \quad &= \quad \Gamma_1\,. \end{aligned} \tag{26}$$

Here $\mathrm{1}_{\Delta_2 *}$ is the push-forward map corresponding to the $(\mathcal{E}_2 : \mathcal{E}_2 * \mathcal{E}_2)$-measurable "diagonal injection" $\mathrm{1}_{\Delta_2} : E_2 \to E_2 \times E_2$. We need to develop the requirements summarized by the last two equations of (26) into specifications for m_2' and m_2'' which are compatible with the two marginalization requirements summarized in the second equation. Note that the marginalization requirements hold at the level of $\mathcal{E}_1$:

$$\Pi^{\pm}{}_* m_2'|_{\mathcal{E}_1 * \mathcal{E}_1} \quad = \quad \hat{\mu}_2|_{\mathcal{E}_1} + \eta_1 + \nu_1^{\pm} \quad = \quad \mu_1^{\pm}\,.$$

Consider $\nu_1^{\pm}$ the marginals of $m_2''|_{\mathcal{E}_1 * \mathcal{E}_1} = \Gamma_1$, the last equation of (26). We wish to extend these $\mathcal{E}_1$ measures into $\mathcal{E}_2$ measures similar to $\nu_2^{\pm}$ in order to connect to the marginals arising in the second equation of (26). Now

$$\begin{aligned} & \frac{(f_1^{\pm} - f_1^{\mp}) \wedge 0}{\int ((f_2^{\pm} - f_2^{\mp}) \wedge 0)\,\mathrm{d}\,p_{1\to 2}}\,\mathrm{d}\,\nu_2^{\pm} \\ = \quad & \frac{(f_1^{\pm}(x_1) - f_1^{\mp}(x_1)) \wedge 0}{\int ((f_2^{\pm}(x_1, y_2) - f_2^{\mp}(x_1, y_2)) \wedge 0) p_{1\to 2}(x_1,\, \mathrm{d}\,y_2)}\,((f_2^{\pm}(x_1, x_2) - f_2^{\mp}(x_1, x_2)) \wedge 0)\,\overline{\mu}_2(\mathrm{d}\,x_1,\, \mathrm{d}\,x_2) \\ = \quad & \frac{(f_2^{\pm}(x_1, x_2) - f_2^{\mp}(x_1, x_2)) \wedge 0}{\int ((f_2^{\pm}(x_1, y_2) - f_2^{\mp}(x_1, y_2)) \wedge 0) p_{1\to 2}(x_1,\, \mathrm{d}\,y_2)}\,p_{1\to 2}(x_1,\, \mathrm{d}\,x_2) \times ((f_1^{\pm}(x_1) - f_1^{\mp}(x_1)) \wedge 0)\,\overline{\mu}_1(\mathrm{d}\,x_1) \\ = \quad & \left(\frac{(f_2^{\pm} - f_2^{\mp}) \wedge 0}{\int ((f_2^{\pm} - f_2^{\mp}) \wedge 0)\,\mathrm{d}\,p_{1\to 2}}\,\mathrm{d}\,p_{1\to 2}\right)\mathrm{d}\,\nu_1^{\pm}, \end{aligned}$$

where the last equality follows from (16). For our purposes, the key point is that the kernel measures given by

$$\lambda^{\pm}(x_1,\, \mathrm{d}\,x_2) \quad = \quad \frac{(f_2^{\pm}(x_1, x_2) - f_2^{\mp}(x_1, x_2)) \wedge 0}{\int ((f_2^{\pm}(x_1, y_2) - f_2^{\mp}(x_1, y_2)) \wedge 0) p_{1\to 2}(x_1,\, \mathrm{d}\,y_2)}\,p_{1\to 2}(x_1,\, \mathrm{d}\,x_2) \tag{27}$$

each have unit total integral when integrated over x_2 and holding x_1 fixed.

Consider η_1 the common marginal of $m_2'|_{\mathcal{E}_1 * \mathcal{E}_1} = \mathrm{1}_{\Delta_2 *}\eta_1$, the penultimate equation of (26). Again, we wish to extend this $\mathcal{E}_1$ measure into an $\mathcal{E}_2$ measure similar to $\nu_2^{\pm}$ in order to connect to the marginals arising in the second equation of (26). Now we can use (23) to deduce

$$
\begin{aligned}
& \left(1 - \frac{(f_1^{\pm} - f_1^{\mp}) \wedge 0}{\int ((f_2^{\pm} - f_2^{\mp}) \wedge 0)\, \mathrm{d}\, p_{1\to 2}}\right) \mathrm{d}\, \nu_2^{\pm} \\
= & \frac{\int (f_2^{\pm} - (f_2^{\pm} \wedge f_2^{\mp}))\, \mathrm{d}\, p_{1\to 2} - (f_1^{\pm} - (f_1^{\pm} \wedge f_1^{\mp}))}{\int ((f_2^{\pm} - f_2^{\mp}) \wedge 0)\, \mathrm{d}\, p_{1\to 2}}\ \mathrm{d}\, \nu_2^{\pm} \\
= & \frac{(f_1^{\pm} \wedge f_1^{\mp}) - \int (f_2^{\pm} \wedge f_2^{\mp})\, \mathrm{d}\, p_{1\to 2}}{\int ((f_2^{\pm} - f_2^{\mp}) \wedge 0)\, \mathrm{d}\, p_{1\to 2}}\ \mathrm{d}\, \nu_2^{\pm} \\
= & \frac{(f_1^{\pm} \wedge f_1^{\mp}) - \int (f_2^{\pm} \wedge f_2^{\mp})\, \mathrm{d}\, p_{1\to 2}}{\int ((f_2^{\pm} - f_2^{\mp}) \wedge 0)\, \mathrm{d}\, p_{1\to 2}}\ (f_2^{\pm} - f_2^{\mp}) \wedge 0)\, \mathrm{d}\, \overline{\mu}_2 \\
= & \frac{(f_2^{\pm} - f_2^{\mp}) \wedge 0)}{\int ((f_2^{\pm} - f_2^{\mp}) \wedge 0)\, \mathrm{d}\, p_{1\to 2}} \left((f_1^{\pm} \wedge f_1^{\mp}) - \int (f_2^{\pm} \wedge f_2^{\mp})\, \mathrm{d}\, p_{1\to 2}\right) \mathrm{d}\, \overline{\mu}_2 \\
= & \left(\frac{(f_2^{\pm} - f_2^{\mp}) \wedge 0)}{\int ((f_2^{\pm} - f_2^{\mp}) \wedge 0)\, \mathrm{d}\, p_{1\to 2}}\ \mathrm{d}\, p_{1\to 2}\right) \left((f_1^{\pm} \wedge f_1^{\mp}) - \int (f_2^{\pm} \wedge f_2^{\mp})\, \mathrm{d}\, p_{1\to 2}\right) \mathrm{d}\, \overline{\mu}_1 \\
= & \lambda^{\pm}(x_1, \mathrm{d}\, x_2)\, \eta_1(\mathrm{d}\, x_1).
\end{aligned}
$$

We can now give the required additive components m_2' and m_2'':

$$
\begin{aligned}
m_2'(\mathrm{d}\, x_1^+, \mathrm{d}\, x_2^+, \mathrm{d}\, x_1^-, \mathrm{d}\, x_2^-) &= (\imath_{\Delta_2 *}\eta_1)(\mathrm{d}\, x_1^+, \mathrm{d}\, x_1^-)\, \lambda^+(x_1^+, \mathrm{d}\, x_2^+)\, \lambda^-(x_1^-, \mathrm{d}\, x_2^-), && (28)\\
m_2''(\mathrm{d}\, x_1^+, \mathrm{d}\, x_2^+, \mathrm{d}\, x_1^-, \mathrm{d}\, x_2^-) &= \Gamma_1(\mathrm{d}\, x_1^+, \mathrm{d}\, x_1^-)\, \lambda^+(x_1^+, \mathrm{d}\, x_2^+)\, \lambda^-(x_1^-, \mathrm{d}\, x_2^-). && (29)
\end{aligned}
$$

By construction, the marginals of $m_2' + m_2''$ are

$$
\left(1 - \frac{(f_1^{\pm} - f_1^{\mp}) \wedge 0}{\int ((f_2^{\pm} - f_2^{\mp}) \wedge 0)\, \mathrm{d}\, p_{1\to 2}}\right) \mathrm{d}\, \nu_2^{\pm} + \frac{(f_1^{\pm} - f_1^{\mp}) \wedge 0}{\int ((f_2^{\pm} - f_2^{\mp}) \wedge 0)\, \mathrm{d}\, p_{1\to 2}}\ \mathrm{d}\, \nu_2^{\pm} \quad = \quad \mathrm{d}\, \nu_2^{\pm},
$$

as required to ensure that $m_2 = \imath_{\Delta_2 *}\hat{\mu}_2 + m_2' + m_2''$ has marginals $\mu_2^{\pm}$, as required by the second line of (26). Inheritance at the level of $\mathcal{E}_1$ follows directly from (28) and (29) because the kernel measures $\lambda^{\pm}$ integrate to 1 as specified after (27). □

Remark 24. As in Lemma 19, copula theory could be used to (over-)parametrize a general family of maximal couplings, by using copulae (parametrized by x_1^+ and x_1^-) to replace the product kernel measures $\lambda^+ \otimes \lambda^-$ in (28) and (29). However, there are varieties of *MEXIT* which are *not* produced in this simple way; in (29) we could replace the parametrized product measure $\lambda^+(x_1^+, \mathrm{d}\, x_2^+)\, \lambda^-(x_1^-, \mathrm{d}\, x_2^-)$ by parametrized measures which are global rather than local copulae: the copula property can be arranged to fail under fixed x_1^+ and x_1^-, even though the correct marginal and inheritance properties are preserved globally.

4.3 Unbounded and/or continuously varying time

MEXIT for all times (with no upper bound on time) follows easily so long as the Kolmogorov Extension Theorem (Doob, 1994, §V.6) can be applied. This is the case, for example, if the state-space is Polish; we state this formally as a corollary to Theorem 22 of the previous section. (For an example of what can go wrong for the Kolmogorov Extension Theorem in a more general measure-theoretic context, see Stoyanov, 1997, §2.3.)

Corollary 25. *Consider two discrete-time random processes X^+ and X^-, begun at the same fixed initial point, taking values in a measurable state-space $(E, \mathcal{E})$* which is a Polish space. MEXIT *couplings exist through all time.*

Under the requirement of Polish state-space, it is also straightforward to establish a continuous-time version of the *MEXIT* result for càdlàg processes. The result requires this preliminary elementary properties about joint laws with given marginals.

Lemma 26. *Suppose that $\{X_i^+\}$ and $\{X_1^-\}$ are two collections of random variables on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ taking values on a metric space (E, d). Suppose that $\{\mathcal{L}\left(X_i^+\right)\}$ and $\{\mathcal{L}\left(X_i^-\right)\}$ are both tight. Then $\{\mathcal{L}\left(X_i^+, X_i^-\right)\}$ is tight on $(E \times E, \tilde{d})$ where $\tilde{d}$ denotes (for instance) the Euclidean product measure $d \times d$.*

Proof. For any $\epsilon > 0$, we can find compact sets S^+, S^- such that $\mathbb{P}(X_i^+ \in S^+) > 1 - \epsilon/2$ and $\mathbb{P}(X_i^- \in S^-) > 1 - \epsilon/2$ for all i. But $S^+ \times S^-$ is $\tilde{d}$–compact and clearly $\mathbb{P}((X_i^+, X_i^-) \in S^+ \times S^-) > 1 - \epsilon$, so that $\{\mathcal{L}\left(X_i^+, X_i^-\right)\}$ is tight on $(E \times E, \tilde{d})$. □

Theorem 27. *Consider two continuous-time random processes X^+ and X^-, begun at the same fixed initial point, with càdlàg paths in a given complete separable metric state-space E (with Borel σ-algebra).* MEXIT *couplings exist through all time.*

Proof. We work first up to a fixed time T.

The space of càdlàg paths in a complete separable metric state-space over a fixed time interval $[0,T]$ can be considered as a Polish space (Maisonneuve, 1972, Théorème 1), using a slight modification of the Skorokhod metric, namely the following *Maisonneuve distance*: if $\tau(t) : [0,T] \to [0,T]$ is a non-decreasing function determining a change of time, and if $|\tau| = \sup_t |\tau(t) - t| + \sup_{s \neq t} \log\left(\frac{\tau(t)-\tau(s)}{t-s}\right)$, then the Maisonneuve distance is given by

$$\mathrm{dist}_M(\omega, \widetilde{\omega}) \quad = \quad \inf_{\tau} \{|\tau| + \mathrm{dist}_E((\omega \circ \tau) - \widetilde{\omega})\}, \tag{30}$$

where ω and $\widetilde{\omega}$ are two càdlàg paths $[0,T] \to E$.

Consider a sequence of discretizations σ_n ($n = 1, 2, \ldots$) of time-space $[0,T]$ whose meshes tend to zero, each discretization being a refinement of its predecessor. Let $X^{\pm,n}(t) = X^{\pm}(\inf\{s \in \sigma_n : s \geq t\})$ define discretized approximations of $X^{\pm}$ with respect to the discretization σ_n. Invoking Theorem 22, we require $X^{+,n}$, $X^{-,n}$ to be maximally coupled as discrete-time random processes sampled only at the discretization σ_n: since they are constant off σ_n, this extends to a maximal coupling of $X^{+,n}$, $X^{-,n}$ viewed as piecewise-constant processes defined over all continuous time.

For a given càdlàg path ω, the discretization of ω by σ_n converges to ω in Maisonneuve distance. This follows by observing that, for each fixed $\varepsilon > 0$, the time interval $[0,T]$ can be covered by pointed open intervals $t \in (t_-, t_+)$ such that $|\omega(s) - \omega(t-)| < \varepsilon/2$ if $s \in (s_-, t)$ and $|\omega(s) - \omega(t)| < \varepsilon/2$ if $s \in (t, s_+)$. By compactness we can select a finite sub-cover. For sufficiently fine discretizations σ we can then ensure the Maisonneuve distance between ω and the resulting discretization is smaller than ε. Consequently, the sequence of joint distributions of $(X^{+,n}, X^{-,n})$ is tight, and we can invoke Lemma 26 to ensure that $\{\mathcal{L}\left(X_i^+, X_i^-\right)\}$ is tight.

Therefore (selecting a weakly convergent subsequence if necessary) we may suppose the joint distribution $(X^{+,n}, Y^{+,n})$ converges weakly to a limit which we denote by $(\widetilde{X}^+, \widetilde{X}^-)$. Note that the marginal laws of $(\widetilde{X}^+, \widetilde{X}^-)$ are the laws of X^+, X^-; however the joint distribution (which is to say, the specific coupling described by $(\widetilde{X}^+, \widetilde{X}^-)$) may well be affected by the exact choice of weakly convergent subsequence.

Suppose time t belongs to one of the discretizations in the sub-sequence, and thus eventually to all (since each discretization is a refinement of its predecessor). Consider the subspace of the Cartesian square of Skorokhod space given by $A_t = [\mathit{MEXIT} \geq t]$. This is a closed subset of the Cartesian square of Skorokhod space. Hence, by the Portmanteau Theorem of weak convergence (Billingsley, 1968, Theorem 2.1),

$$\limsup_{n \to \infty} \mathbb{P}\left[(X^{+,n}, X^{-,n}) \in A_t\right] \quad \leq \quad \mathbb{P}\left[(\widetilde{X}^+, \widetilde{X}^-) \in A_t\right].$$

If n is large enough so that $t \in \sigma_n$, then $\mathbb{P}\left[(X^{+,n}, X^{-,n}) \in A_t\right]$ is equal to the probability that *MEXIT* happens at t or later. By maximality of *MEXIT* for the discretized processes at time $t \in \sigma_n$, it follows for such n that $\mathbb{P}\left[(\widetilde{X}^+, \widetilde{X}^-) \in A_t\right] = \mathbb{P}\left[(X^{+,n}, Y^{+,n}) \in A_t\right]$ is itself maximal. The càdlàg property then implies maximality of the limiting coupling for all times $t \leq T$. The result on *MEXIT* for all time follows by application of the Kolmogorov Extension theorem as in Corollary 25. This concludes the proof. □

Remark 28. Sverchkov and Smirnov (1990) prove a similar result for maximal couplings by means of general martingale theory.

Remark 29. Note that Théorème 1 of Maisonneuve (1972) can be viewed as justifying the notion of the space of càdlàg paths: this space is the completion of the space of step functions under the Maisonneuve distance dist_M. Thus in some sense Theorem 27 is the maximal reasonable result concerning *MEXIT*!

5 *MEXIT* for diffusions

The results of Section 4 apply directly to diffusions, which therefore exhibit *MEXIT*. This section discusses the solution of a *MEXIT* problem for Brownian motions, which can be viewed as the limiting case for random walk *MEXIT* problems.

It is straightforward to show that *MEXIT* will generally have to involve constructions not adapted to the shared filtration of the two diffusion in question. By "faithful" *MEXIT* we mean a *MEXIT* construction which

generates a coupling between the diffusions which is faithful / immersed with respect to the joint and individual filtrations. We consider the case of regular diffusions X^+ and X^- (strictly elliptic, with smooth coefficients).

Theorem 30. *Suppose X^+ and X^- are coupled regular diffusions, thus with continuous semimartingale characteristics given by their drift vectors and volatility (infinitesimal quadratic variation) matrices, begun at the same point, with this initial point lying in the open set where either or both of the drift and volatility characteristics disagree. Faithful* MEXIT *must happen immediately.*

Proof. Let T be the *MEXIT* time, which by faithfulness will be a stopping time with respect to the common filtration. If X^+ and X^- are semimartingales agreeing up to the random time T, then the localization theorems of stochastic calculus tell us that the integrated drifts and quadratic variations of X^+ and X^- must also agree up to time T. It follows that X^+ and X^- agree as diffusions up to time T. Were the faithful *MEXIT* stopping time to have positive chance of being positive then the diffusions would have to agree on the range of the common diffusion up to faithful *MEXIT*; this would contradict our assertion that the initial point lies in the open set where either or both of the drift and volatility characteristics disagree. □

By way of contrast, *MEXIT* can be described explicitly in the case of two real Brownian motions X^+ and X^- with constant but differing drifts. Because of re-scaling arguments in time and space, there is no loss of generality in supposing that both X^+ and X^- begin at 0, with X^+ having drift $+1$ and X^- having drift -1.

Theorem 31. *If $X^\pm$ is Brownian motion begun at 0 with drift ± 1, then* MEXIT *between X^+ and X^- exists and is almost surely positive.*

Proof (outline). The proof is derived in Subsection 5.2 below as a limiting version of the random walk argument in Subsection 3.2.1. Alternatively one can argue succinctly and directly using the excursion-theoretic arguments of Williams' (1974) celebrated path-decomposition of Brownian motion with constant drift (an exposition in book form is given in Rogers and Williams, 1987).

Calculation shows that the *bounded* positive excursions of X^+ (respectively $-X^-$) from 0 are those of the positive excursions of a Brownian motion of *negative* drift -1, while the *bounded* negative excursions of X^+ (respectively $-X^-$) from 0 are those of the negative excursions of a Brownian motion of *positive* drift -1. (The unbounded excursion of X^+ follows the law of the distance from its starting point of Brownian motion in hyperbolic 3-space, while the unbounded excursion of X^- has the distribution of the mirror image of the unbounded excursion of X^+.)

Viewing $X^\pm$ as generated by Poisson point processes of excursions indexed by local time, it follows that we may couple X^+ and X^- to share the same bounded excursions, with unbounded excursions being the reflection of each other in 0. Moreover the processes have disjoint support once they become different. So the Recognition Lemma for Maximal Coupling (Lemma 21) applies, and hence this is a *MEXIT* coupling. □

5.1 Explicit calculations for Brownian *MEXIT*

Let X^+ and X^- begin at 0, with X^+ having drift $+\theta$ and X^- having drift $-\theta$ with $\theta > 0$. The purpose of this section is to offer explicit calculations of *MEXIT* and *MEXIT* means.

Calculation 1. The meet of the distributions of X_t^+ and X_t^- is the meet of $N(\theta t, t)$ and $N(-\theta t, t)$, and the probability of *MEXIT* happening after time t is given by the total mass of this meet sub-probability distribution. Therefore:

$$\begin{aligned}\mathbb{P}\left[MEXIT \geq t\right] &= \mathbb{P}\left[N(0,t) < -\theta t\right] + \mathbb{P}\left[N(0,t) > \theta t\right] \\ &= 2\,\mathbb{P}\left[N(0,t) > \theta t\right] \\ &= \frac{2}{\sqrt{2\pi}} \int_{\theta\sqrt{t}}^{\infty} e^{-u^2/2} du.\end{aligned}$$

Thus,

$$\mathbb{E}\left[MEXIT\right] = \frac{2}{\sqrt{2\pi}} \int_0^{\infty} \int_{\theta\sqrt{t}}^{\infty} e^{-u^2/2} du dt = \theta^{-2}.$$

Remark 32. Excursion theoretic arguments can be used to confirm this is mean time to *MEXIT* for the specific construction given in Theorem 31.

Calculation 2. We now consider the expected amount of time T during which processes agree before *MEXIT* happens.

$$\begin{aligned}
\mathbb{E}[T] &= \int_0^\infty \mathbb{E}_W\left[\min\{e^{\theta W_t-\theta^2 t/2}, e^{-\theta W_t-\theta^2 t/2}\}\right] dt \\
&= 2\int_0^\infty \mathbb{E}_W\left[e^{-\theta W_t-\theta^2 t/2}; W_t > 0\right] dt \\
&= 2\int_0^\infty \int_0^\infty \frac{1}{\sqrt{2\pi t}} \exp\left(-\frac{(w+\theta t)^2}{2t}\right) dw dt \\
&= \theta^{-2}.
\end{aligned}$$

5.2 An explicit construction for *MEXIT* for Brownian motions with drift

In this section, we continue the scenario of Calculation 2 above. We see that *MEXIT* should have the complementary cumulative distribution function

$$\mathbb{P}[MEXIT \geq t] = 2\Phi(-\theta\sqrt{t}), \tag{31}$$

where $\Phi(y) = \int_{-\infty}^{y}(2\pi)^{-1/2}e^{-u^2/2}du$. A natural question to ask is as follows: how can one explicitly construct and understand this *MEXIT* time in a way that relates to the random walk constructions of Subsection 3.2.1? In this section we first deduce a candidate coupling and EXIT time, and then we proceed to show by direct calculation that our construction indeed gives the correct *MEXIT* time distribution above.

We note from the discrete constructions of Section 3 (in particular Subsection 3.2) that *MEXIT* is only possible when the Radon-Nikodym derivative between the "p" and "q" processes moves from being below 1 to above 1 or moves from being above 1 to below 1. Let $\mathbb{P}^+$, $\mathbb{P}^-$ denote the probability laws of X^+, X^- respectively. We have that

$$\frac{d\mathbb{P}^+}{d\mathbb{P}^-}(W_{[0,T]}) = \exp\{2\theta W_T\},$$

which is continuous in time with probability 1 under both $\mathbb{P}^+$ and $\mathbb{P}^-$. By analogy to the discrete case, the region in which *MEXIT* could possibly occur corresponds to the interface $\frac{d\mathbb{P}^+}{d\mathbb{P}^-}(W_{[0,T]}) = 1$ (that is, where $W_T = 0$).

Now we shall focus on the random walk example at the end of Subsection 3.2. We note that the *MEXIT* distribution given in (11) can be constructed as the first time the occupation time of 0 exceeds a geometric random variable with "success" probability $1-2\eta$. We aim to give a similar interpretation for the Brownian motion case. To do this, we shall use a sequence of random walks converging to the appropriate Brownian motions. To this end, let

$$\eta_n = \frac{1}{2}\left(1-\frac{\theta}{n}\right),$$

and set X^{n+} and X^{n-} to be the respective simple random walks with up probability $1-\eta_n$ and η_n and sped up by factor n^2. We assume (unless otherwise stated) that all processes begin at 0 so that we have that

$$X^{n+}(t) = \sum_{i=1}^{\lfloor n^2 t\rfloor} Z_i^{n+},$$

where $\{X_i^{n+}\}$ denote dichotomous random variable taking the value $+1$ with probability $1-\eta_n$ and -1 with probability η_n. We define X^{n-} analogously.

Given this setup, we have the classical weak convergence results that the law of X^{n+} converges weakly to that of X^+, and similarly X^{n-} converges weakly to X^-. Moreover the joint pre-*MEXIT* process described in Subsection 3.2 will have drift $-sgn(X_t)\theta$. The following holds for the *MEXIT* probability in (11)

$$\mathbb{P}[MEXIT > t] = \left(1-\frac{\theta}{n}\right)^{n\ell_t^n} \longrightarrow e^{-\theta\ell_t^n},$$

where ℓ_t^n is the Local Time at 0 of the pre-*MEXIT* process for the nth approximation random walk.

In the (formal) limit as $n\to\infty$, this recovers the construction in Theorem 31 of Brownian motion *MEXIT* time, as follows. Let X be the diffusion with drift $-\operatorname{sgn}(X)\theta$ and unit diffusion coefficient started at 0 and let

ℓ_t denote its local time at level 0 and time t. Then set E to be an exponential random variable with mean θ^{-1}. Then the pre-*MEXIT* dynamics are described by X until $\ell_t > E$ at which time *MEXIT* occurs.

We shall now verify that this construction does indeed achieve the valid *MEXIT* probability given in (31). By integrating out E we are required to show that

$$\mathbb{E}\left[e^{-\theta \ell_t}\right] = 2\Phi(-\theta\sqrt{t}) \ .$$

We proceed to do so. Firstly, we note that by symmetry, we may set ℓ_t to be the local time at level 0 of Brownian motion with drift $-\theta$ reflected at 0. Note that by an extension of Lévy's Theorem (see Peskir, 2006) that the law of ℓ_t is the same as that of the maximum of Brownian motion with drift θ, i.e. that of X^+. Now this law is well-known as the Bachelier-Lévy formula (see for example Lerche, 2013):

$$\mathbb{P}\left[\ell_t < a\right] = \Phi\left(\frac{a}{\sqrt{t}} - \theta\sqrt{t}\right) - e^{2a\theta}\Phi\left(\frac{-a}{\sqrt{t}} - \theta\sqrt{t}\right),$$

with density

$$f_{\ell_t}(a) = \frac{1}{\sqrt{t}}\left(\phi\left(\frac{a}{\sqrt{t}} - \theta\sqrt{t}\right) + e^{2a\theta}\phi\left(\frac{-a}{\sqrt{t}} - \theta\sqrt{t}\right) - 2\sqrt{t}\theta e^{2a\theta}\Phi\left(\frac{-a}{\sqrt{t}} - \theta\sqrt{t}\right)\right),$$

where ϕ is the standard normal density function $\phi(y) = (2\pi)^{-1/2}e^{-y^2/2}$. By direct manipulation of the exponential quadratic in the second of the three terms above, it can readily be shown to equal the first term. Thus

$$f_{\ell_t}(a) = \frac{1}{\sqrt{t}}\left(2\phi\left(\frac{a}{\sqrt{t}} - \theta\sqrt{t}\right) - 2\sqrt{t}\theta e^{2a\theta}\Phi\left(\frac{-a}{\sqrt{t}} - \theta\sqrt{t}\right)\right).$$

We now directly calculate the Laplace transform of this distribution to obtain (31).

$$\begin{aligned}\mathbb{E}\left[e^{-\theta \ell_t}\right] &= \frac{2}{\sqrt{t}}\int_0^\infty e^{-\theta a}\left(\phi\left(\frac{a}{\sqrt{t}} - \theta\sqrt{t}\right) - \sqrt{t}\theta e^{2a\theta}\Phi\left(\frac{-a}{\sqrt{t}} - \theta\sqrt{t}\right)\right) da \\ &=: \frac{2}{\sqrt{t}}(T_1 - T_2) \ .\end{aligned}$$

Using integration by parts, we easily work with T_2 to obtain

$$T_2 = T_1 - \sqrt{t}\Phi(-\theta\sqrt{t}),$$

which implies the assertion in (31), as required.

6 Conclusion

In this paper, we have studied an alternative coupling framework in which one seeks to arrange for two different Markov processes to remain equal for as long as possible, when started in the same state. We call this "un-coupling" or "maximal agreement" construction MEXIT, standing for "maximal exit" time. *MEXIT* sharply differs from the more traditional maximal coupling constructions studied in Griffeath (1975), Pitman (1976), and Goldstein (1979) in which one seeks to build two different copies of the same Markov process started at two different initial states in such a way that they become equal as soon as possible.

This work begins with practical motivation for *MEXIT* by highlighting the importance of un-coupling/maximal agreement arguments in a few key statistical and probabilistic settings. With this motivation established, we develop an explicit *MEXIT* construction for Markov chains in discrete time with countable state-space. We then generalize the construction of *MEXIT* to random processes on general state-space in continuous time. We conclude with the solution of a *MEXIT* problem for Brownian motions.

As noted in Remark 8, the approach that we have followed in the construction of *MEXIT* introduces the role of copula theory in parametrising varieties of maximal couplings for random processes. Our future work will aim to establish a definitive role for *MEXIT* (as well as for probabilistic coupling theory in general) in copula theory.

Acknowledgements The first author thanks Professor Larry Shepp and Quan Zhou for helpful conversations. The second author was supported by EPSRC grant EP/K013939. The third author was supported by EPSRC grants EP/K034154/1, EP/K014463/1, EP/D002060/1. The fourth author received research support from NSERC of Canada.

References

Aldous, D. (1983). Random walks on finite groups and rapidly mixing Markov chains. In *Séminaire de Probabilités XVII 1981/82*, pp. 243–297. Springer.

Alquier, P., N. Friel, R. Everitt, and A. Boland (2016). Noisy Monte Carlo: onvergence of Markov chains with approximate transition kernels. *Statistics and Computing 26*(1-2), 29–47.

Andrieu, C. and G. O. Roberts (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics 37*(2), 697–725.

Billingsley, P. (1968). *Convergence of probability measures*. New York: John Wiley & Sons Inc.

Burdzy, K. and W. S. Kendall (2000, May). Efficient Markovian couplings: examples and counterexamples. *The Annals of Applied Probability 10*(2), 362–409.

Ceperley, D. M. and M. Dewing (1999). The penalty method for random walks with uncertain energies. *The Journal of Chemical Physics 110*(20), 9812–9820.

Dellacherie, C. and P.-A. Meyer (1979, January). Probabilities and potential. A.

Doob, J. L. (1994). *Measure Theory*, Volume 143 of *Graduate Texts in Mathematics*. New York, NY: Springer New York.

Farrell, R. H. (1964). Asymptotic Behavior of Expected Sample Size in Certain One Sided Tests. *The Annals of Mathematical Statistics 35*(1), 36–72.

Goldstein, S. (1979). Maximal coupling. *Probability Theory and Related Fields 46*(2), 193–204.

Griffeath, D. (1975). A maximal coupling for Markov chains. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete 31*, 95–106.

Halmos, P. R. (1978, February). *Measure Theory* (1st ed. 19 ed.). New York: Springer.

Kendall, W. S. (2015). Coupling, local times, immersions. *Bernoulli 21*(2), 1014–1046.

Lerche, H. R. (2013). *Boundary crossing of Brownian motion: Its relation to the law of the iterated logarithm and to sequential analysis*, Volume 40. Springer Science & Business Media.

Lindvall, T. (1992). *Lectures on the coupling method*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. New York: John Wiley & Sons Inc.

Maisonneuve, B. (1972). Topologies du type de Skorohod. *Séminaire de probabilités de Strasbourg 6*, 113–117.

Medina-Aguayo, F., A. Lee, and G. O. Roberts (2015). Stability of noisy Metropolis–Hastings. *Statistics and Computing*, 1–25.

Nelson, R. B. (2006). *An introduction to copulas*. Springer Series in Statistics. New York, NY: Springer New York.

Nicholls, G. K., C. Fox, and A. M. Watt (2012). Coupled MCMC with a randomized acceptance probability. *arXiv preprint arXiv:1205.6857*.

Peskir, G. (2006). On reflecting brownian motion with drift. In *Proc. Symp. Stoch. Syst.*, pp. 1–5.

Pitman, J. W. (1976). On coupling of Markov chains. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete 35*(4), 315–322.

Roberts, G. O., A. Gelman, and W. R. Gilks (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms. *The Annals of Applied Probability 7*(1), 110–120.

Roberts, G. O. and J. S. Rosenthal (2004). General state space Markov chains and MCMC algorithms. *Probability Surveys 1*, 20—-71 (electronic).

Roberts, G. O. and J. S. Rosenthal (2007). Coupling and Ergodicity of Adaptive MCMC. *Journal of Applied Probability 44*(2), 458–475.

Rogers, L. C. G. and D. Williams (1987). *Diffusions, Markov processes, and martingales*, Volume II. Chichester / New York: John Wiley & Sons Ltd.

Rosenthal, J. S. (1997). Faithful Couplings of Markov Chains: Now Equals Forever. *Advances in Applied Mathematics 18*(3), 372–381.

Stoyanov, J. (1997). *Counterexamples in probability*. Chichester / New York: John Wiley & Sons Ltd.

Sverchkov, M. Y. and S. N. Smirnov (1990). Maximal coupling for processes in D[0,∞]. *Dokl. Akad. Nauk SSSR 311*(5), 1059–1061.

Thorisson, H. (2000). *Coupling, stationarity, and regeneration*. New York: Springer-Verlag.

Völlering, F. (2016, August). On maximal agreement couplings. *arXiv 1608.01511*, 10pp.

Williams, D. (1974, June). Path decomposition and continuity of local time for one-dimensional diffusions, I. *Proceedings of the London Mathematical Society s3-28*(4), 738–768.